COMPUTER FACILITIES OPTIMIZATION STUDIES

Final Technical Report

NASA-Cooperative Agreement NCC2-441

for the period
September 1, 1986 - December 31, 1992

Submitted to

National Aeronautics and Space Administration
Ames Research Center
Moffett Field, California 94035

Computer Systems and Research Division
Marcelline C. Smith
Chief and Technical Officer

Prepared by

ELORET INSTITUTE
1178 Maraschino Drive
Sunnyvale, CA 94087
Phone: 408 730-8422 and 415 493-4710
Telefax: 408 730-1441
K. Heinemann, President and Grant Administrator
Gilbert C. Lyle, Principal Investigator
S. Strauss-Duckett, Co-Principal Investigator
Terry C. Nelson, Co-Principal Investigator

4 March, 1993

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.

4. Title and Subtitle	5. Report Date
Computer Facilities Optimization Studies	4 March 1993
	6. Performing Organization Code
7. Author(s)	8. Performing Organization Report No.
Gilbert C. Lyle	
9. Performing Organization Name and Address	10. Work Unit No.
Eloret Institute 3788 Fabian Way Palo Alto CA 94303	11. Contract or Grant No. NCC2-441
12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
National Aeronautics and Space Administration, Washington, D.C. 20456	09/01/86 - 12/31/92
	14. Sponsoring Agency Code

15. Supplementary Notes

Point of Contact: Mr. Gilbert C. Lyle
c/o 233-10, NASA Ames Research Center, Moffett Field, CA 94035

16. Abstract

A comprehensive plan was developed to determine the best options for integrating and upgrading various ADP systems in scientific research environments, such as at NASA Ames Research Center.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
ADP Workstations Networks	Unclassified, Unlimited

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price*
Unclassified	Unclassified		

*For sale by the National Technical Information Service, Springfield, Virginia 22161

NASA CONTRACTOR REPORT

COMPUTER FACILITIES OPTIMIZATION STUDIES

Gilbert C. Lyle

Eloret Institute
3788 Fabian Way
Palo Alto, CA 94303

Prepared for

Ames Research Center
under **Cooperative Agreement NCC2-441**

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035

NASA CONTRACTOR REPORT

STATISTICAL DISCRETE PARTICLE SIMULATION STUDIES

Gilbert C. Lyle

CONTRACT NAS2–

NASA

II

September, 1986 through August 1987

At its creation in 1986 the Thermosciences Division acquired various research tools including computer terminals, a gateway size machine, various facilities for taking physical data, and the computers ancillary to these research facilities. Some of that equipment was relatively new, some was aging rapidly, and some was so outdated that it imposed pre-Apollo work methods on the solutions of then current problems.

The Principal Investigator proposed, and management agreed, that a comprehensive plan should be developed to determine the best options for integrating the ADP systems available at each level of the Division and for upgrading where necessary. It was recognized that if the system were to continue to be developed to maximize effectiveness, in terms of the researchers, rather than efficiency, in terms of the equipment, then clear and realistic priorities would need to be developed.

As a result, this Cooperative Agreement was entered into to conduct studies that would result in recommendations to Branch and Division management for the setting of priorities and in incorporating them into an overall ADP plan for the organization. That plan was to enable Code RT management to provide both scientists and managerial staff with the computer facilities that they would require to solve the problems of the upcoming decade.

Early in the life of the Agreement an analysis of the options available for upgrading the Division's principal local computer resource was prepared and presented to Division management.[1] These recommendations resulted in several major upgrades. Prominent among them was the installation of four more megabytes of memory into the VAX 11/785, upgrading its operating system to VMS 4.3, and the addition of a QMS 2400 Laser Printer.

In order to carry out its assignment of adding real chemistry effects to the computational fluid dynamics codes that were to be used to design the National AeroSpace Plane the Division needed to increase its computer resources. Studies by the Principal Investigator led both Branch and Division management to conclude that it was not practical to simply upgrade the local distributed computer facilities, which were centered around the DEC VAX 11/785, "H. JULIAN ALLEN." Simply stated, the three arguments against such an upgrade were:

1. HJA could not remain in its then current location (the 3.5 ft. Wind Tunnel control room) after the Tunnel became fully operational;
2. The Division lacked a suitable space in which to create a computer site;
3. The Division could not afford the time that would be required to get a major upgrade through the ADP procurement system.

In discussions with the Technical Officer of this Cooperative Agreement it was agreed that the Principal Investigator would help the Division to develop an ADP plan that would enable them to meet their immediate goals and to allow them to get out of the

distributed computer business with a minimum of disruption. It was expected that the size of the user group, particularly the computational fluid dynamists in Code RTA, would increase substantially so the Principal Investigator's plan recommended an approach that would split the user community into those who would remain on HJA and those who would move to a facility provided by the CCF. The group who were to move to the new system, perhaps some 25 persons, was to be comprised of the computational fluid dynamists and the most productive of the computational chemists. This placed a requirement on the new facility that it be reliable from the beginning.

The new system was to be used to edit code, pre- and post-process CFD output files, and serve as a gateway to the supercomputers in much the same way that HJA currently served the Division. Therefore, it had to have access to the CRAY XMP, the NAS CRAY, the CYBER 205, and the MASS STORE. We expected to transfer post-processed CFD output to IRIS workstations, located in Bldg. 230, for analysis. The hard copy output was generally to be sent to printers located at the node STS.

The two branches, RTA and RTC, agreed that providing an upgrade to their computer capabilities was an activity of the highest priority and so they would fund this effort with up to $250,000 apiece. Further, they agreed to solicit money from NASA Headquarters specifically for this purpose. If additional funds became available, the Division would be able to match their funding up to $500,000.

The principal investigator negotiated with Code RCE for the return of Codes RTA and RTC to a machine in the Central Facility VAX farm. When Code RTA and the Division withdrew because of funding problems the negotiations were continued and ultimately lead to the assignment of a machine for the use of the computational chemists alone.

One of the main reasons that users originally moved from the CCF VAX cluster to distributed minicomputers was their inability to effectively manage their share of the centralized resource. Both the management of Code RCE and of Code RTC were in agreement that they must jointly develop a workable scheme to solve this problem. Therefore, the Principal Investigator drew up a plan to allocate management duties between the central system and the local system manager for the new resource (one half of "JUPITER", changing to all of "JUPITER" in September 1987). Briefly, the responsibilities of the Branch would be:

1. To provide a central point for handling user requests or problems;
2. To bring any plans to modify the system to the attention of the users in order to determine if the changes and the timing were in their best interests;
3. To work closely with central system management to see that any system work, including maintenance and enhancements, was done with the least amount of inconvenience to the users;
4. To report any performance degradation detected from monitoring the system or by the user community to CCF staff;
5. To see that users had as much advance notice as possible of scheduled downtimes, and to be aware of special projects that were time critical;
6. To identify alternatives for resources that were not presently available;

7. To develop system software as requested by users.

In the Spring of 1987 Ames management, in the person of Deputy Director Dale Compton, solicited the views of the Branch Chief of the Computational Chemistry Branch on the possibilities of establishing the position of ADP Manager in each of the Center's Divisions. In discussions with the Technical Officer of the Agreement it was decided that the Principal Investigator would provide an analysis of that proposal.[2]

The state of the ADP planning effort within the Computational Chemistry Branch and the Thermosciences Division at the end of the Agreement year are fully covered in the year-end report dated 1 September 1987 and so will not be repeated here.

September 1987 through August 1988

Over and above the events reported in the mid-year and year-end reports the Principal Investigator was involved in the following activities.

- Devised and negotiated a joint Code RC-RT pilot project which examined productivity increases made possible by providing the RTC scientists with workstations.
- Recommended a complete re-evaluation and re-writing of the Support Services contract in the light of the significant changes in the ADP resources of the Computational Chemistry Branch.
- Devised and recommended a new management scheme for the "H Julian Allen" VAX which took into account the major changes in the Divisions ADP resources.
- An investigation was undertaken to determine the amount and type of work being done by the Code RTC computational chemists during non-duty hours. Various accounting procedures were examined and combined and then used to gather usage data for the Branch's VAX, "JUPITER". These data were analyzed and used to determine the hardware and software requirements for the proposed Scientist's Workstation.[3]
- Planned, and prepared documentation for the acquisition of a Code RT/RFE Output Station, a group of Code RFE Scientific Document Workstations, Code RTC Scientist's Workstations and Code RT Administrative Workstations.

September 1988 through August 1989

In October of 1988, as the result of discussions with the Technical Officer, the Principal Investigator developed a position paper on the subject of Center management assigning priorities for the use of computer resources at the Ames.[4]

The majority of the year was spent in aiding the Division and its Branches in the implementation of the ADP plans that the Principal Investigator had developed during the previous year. The organization now managed three VAXs ("H JULIAN ALLEN",

"JUPITER" and "CALLISTO"), two MICROVAXs ("MARV" and "MOE") and five major printers, both line and laser. They were also responsible for the following, existing and planned, communications facilities:

1. "H JULIAN ALLEN" Micom Switch;
2. Code RTC 470 Instamux - Micom network;
3. Code RTC Ethernet network;
4. Code RTC Proteon network;
5. Code RT Administrative PC network.

The Principal Investigator developed programs to collect and analyze data and formulated, alone and in cooperation with persons from Code ED, many test and monitoring procedures which allowed for better defining the Division's requirements for the replacement of the outdated Sytek equipment that had formed the backbone of the Division's computer communications network for so many years. This effort included providing separate solutions for each of Codes RTC, RTM and the rest of the user community.

The Thermal Protection Materials Branch (Code RTM) transferred all of their computer operations from the Division owned VAX 11/785, "H Julian Allen", to "MOE", a Branch designed, installed and owned VAX Cluster. Planning for a second transition to a new system for The Computational Chemistry Branch from their VAX 11/785, Jupiter, to the Center's first Convex 210 was also undertaken.

September 1989 through August 1990

A continuing activity that persisted throughout the year was the development of a strategy that would enable the Division to increase its share of supercomputer resources at the Center. The effort began with an analysis of the Division's use of the CCF supercomputer with respect to total Code R use.[5] The final analysis and recommendations of this effort were offered as this Cooperative Agreement's year-end report

Another major event of the year, covered in the semiannual report, centered around the acquisition, by Code RTC, of a Convex 210.

It was during this year that the Technical Officer of this Cooperative Agreement, who was also the Branch Chief of the Computational Chemistry Branch, left the organization to assume the position of Division Chief of the Numerical Aerodynamic Simulation Systems Division. The role of Technical Officer shifted to James Arnold, the Division Chief of the Thermosciences Division.

September 1990 through August 1991

There was one ongoing project that continued from the first year of the Cooperative

Agreement until the fall of 1991. That effort revolved around the question of whether the NEMS database of the Thermosciences Division's ADP inventory could be a useful tool in the Center's annual ADP Planning effort. The results and conclusions of that study were documented in the mid-year report.

The Division, in a continuation of their efforts to increase access to supercomputers, made use of an analysis, by the Principal Investigator, of the year's allotment of NAS resources.[6]

The original goals of this Cooperative Agreement were judged to have been met by the end of this year and so a final summary paper on the Division's ADP planning capabilities was offered as the year end report.

September 1991 through December 1992

With the termination of the work done in the Thermosciences Division the Cooperative Agreement was moved to the Computer Systems and Research Division with Marcelline C. Smith as Technical Officer. The report for this performance period is attached as Footnote 7.

Footnote 1

CENTER PLAN

The center, under the leadership of Marcie Smith, is about to implement its first comprehensive, coordinated plan to acquire ADP equipment. Under this plan the center will be able to purchase self-consistent, compatible systems over a broad range of capabilities by dealing with a single vendor.

This should result in the development of a standardized system of computer resources throughout the Center without the problems of the past. Up to now the government's policies toward ADP procurements, particularly with respect to sole source justifications, has conflicted directly with the user community's requirement for compatibility with existing systems and network nodes.

As a part of this planning exercise Code RC, again at Marcie's direction, is re-examining its capacity to provide the user community with services. Access to both hardware and software are being looked at in an effort to find ways to provide researchers with the capabilities that they require without burdening their organizations with the costs of maintaining totally separate facilities.

It is in the Division's best interest to take a vigorous role in support of these activities. By becoming actively involved early, Division management will have access to the information it will require to best take advantage of the time lag between the initial phase and the implementation of these plans, to reevaluate its position with relation to its computer resources. It will be recalled that the reason that the Division developed its own distributed computer system was because the shared resource concept was not properly managed at the mini-computer level under the previous CCF/VAX farm arrangement. The new planning effort presents the Division with an ideal opportunity to leap frog the bottlenecks present in its currently saturated computer resource and examine ways to expand to the capabilities required with a minimum of cost in money and time.

HJA UPGRADE

As the possibility of the Division getting funds for an upgrade to its computing facilities waxes and wanes it seems to me to be useful to review the options open to management in the light of the above mentioned changes in the center's ADP situation.

IF THE DIVISION GETS THE MONEY TO UPGRADE

If the funds do indeed become available the Division can replace or augment the VAX 11\785, that makes up HJA, through the normal ADP purchasing process, or it can keep HJA as it is and attempt to provide added capacity through some other mechanism. Upgrading the VAX 11/780 presents three problems of major proportions.

It will be nearly impossible for the Division to replace the VAX within a time frame which will allow it to meet its research deadlines because of the Government's current ADP purchasing requirements. The machine can not be purchased under sole source procedures, so great care will have to be taken at every step to insure compatibility with the existing facilities.

Division policy has been, since it was decided to reactivate the 3.5 Foot Wind Tunnel, to provide a new site for the Division's computer facilities. Site preparation costs for a new location are $150 per square foot, which will significantly increase the capital outlay. In any case, no site large enough to house both the present equipment and any upgrades has been found. Any increase in capability, whether in place of or in addition to, will require a significant increase in operations and maintenance costs. Given the current restrictions that NASA Headquarters has placed on ADP acquisitions it is not at all clear that the

Division would be allowed to upgrade its computer resources if it did have the money.

A better solution, if the Division does get the funds, would seem to be to attempt to purchase the required services from the Central Computer Facility. What Code RT needs is not a VAX 86xx, but the capability for a given number of users to pre-process, submit to, and post-process programs and data from the CCF and NAS supercomputers. Given the choice, the Division would be best off if it did not have to house this capability in its own buildings. It should be remembered that the decision to embrace the concept of distributed processing to the point of buying a computer system stemmed not from a technical but a management problem. The way to address that management problem is to avoid the sharing of, less than supercomputer sized, resources at a wider than Division level. Marcie Smith is by far the most cooperative and research oriented chief the Computer Division has ever had, so it seems very likely that some mutually satisfactory arrangement could be worked out.

IF IT DOESN'T GET THE MONEY

If the Division is required to get through the next fiscal year with no additional money available to upgrade its computer resources it must provide the minimum system that will allow it to function in the most effective way and look to the future, and the Center ADP plan, for any long term improvements.

UPGRADE THE DISK FARM

The most pressing need for increased capabilities that is within the Division's ability to fund is for added disk capacity for HJA. The storage requirements that the new researchers added to Code RTA will have can not be met with the presently available disk space. We can add two disks to the system without adding another disk controller but this will fall short of meeting the expanded requirement. Since the system, as presently configured, will not accept any more disk controllers we will be required to change the way the system interfaces its disk drives.

We are at the limit of our ability to back the system to tape with the present disk capacity and tape drive. We will, therefore, have to upgrade our tape drive at the same time as we increase our disk farm. Unfortunately the bus on the VAX 11\785 is not compatible with the bus on the newest generation of VAXs, so this new tape drive will not be of use if we later upgrade our CPU.

REDUCE THE LOAD ON H.J.A

Increasing the size of the disk farm does not add to our current capabilities, it only allows us to accommodate an expanded group of researchers. As the system is currently saturated to the point of limiting the output of the Division's most productive scientists, expanding the disk farm will cause more problems than it solves if no other measures are taken. The alternative to expanding the system to correspond to usage is to reduce the usage to correspond to system capacity. Management can control access to its computer facilities in a number of ways including assigned time slots, assigned priorities according to projects or individual scientists, and the reduction of non-vital work. This last reduction might be accomplished by putting management personnel on Personal Computers and by cutting back severely on computer access by students and/or visiting professors and scientists.

WORKSTATIONS

At present the Division has no policy, and therefore, no consistency in the sub-VAX systems that it uses. The machines available include IBM PC\XT, IBM compatible AT, non IBM compatible DEC PRO and DECMATE, Apple, and IRIS workstations. Many of the Division's workstations were acquired as gifts, often because no one else wanted them. An expedient way to get equipment, perhaps, but one that hardly makes for an optimum solution of any computing situation. Management needs to decide who, if anyone, is to be shifted to workstations and how those workstations are to be chosen and paid for.

RIACS PROPOSAL

RIACS has suggested that it conduct a comparative study of a workstation/graphics/file server system suitable for use by scientists. If the Division chooses to participate in this study it is possible that not only the long term scientific workstation question might be answered, it is also possible that some scientists might be shifted off of "H JULIAN ALLEN" very soon. It is suggested that this effort be discussed with the appropriate persons in Code RC and RIACS.

WORKSTATIONS FOR AI

Within the next year or so there will be a large number of 32 bit desktop computers on the market and, with the advent of these machines, a significant increase in the availability of software for the development of artificial intelligence codes. The Division's requirement will, however, predate this effort and so make these advances too late to be of much value to us. At the present time the best choice for a small computer capable of working in the field of AI would be one of the desktop sized VAXs. AI development software already exists for them and they are compatible with "HJA".

Footnote 2

Memo

To: David M. Cooper, Chief, Computational Chemistry Branch

From: Gilbert C. Lyle, Eloret Institute

Date: 8 APRIL 1987

Subject: Proposed ADP Manager position

There are three possible approaches to addressing the question of writing a job description for Division level ADP managers. first, we can consider it as a serious attempt to address one of Ames management problems. Second, as a mechanism that Code RT can use to insure control of its share of computer resources. And third, as busy-work to be done to get Dale off our backs. Despite many painful, frustrating interactions with Center Management in the past I think we must, for our own satisfaction, try the first. If we are unable to accomplish the first then we must, for our own safety, try the second. If we are unable to accomplish the second then we must, for our own self-respect, get out of the business.

There are two questions of importance that must be answered if we are to proceed as though this were an attempt to accomplish something of value. The first problem that we are going to have to examine is: How are we going to get anyone else to take this exercise seriously? This may, at first glance, seem to be a ridiculous query. If there is any one at Ames who deserves to be taken seriously it is Dale. But consider, for a moment, what we are up against. Since Sy left, senior management has established a style of decision making based on the formation of non-productive committees and on a policy of disinterest shattered only by occasional intervention based on favoritism. Not surprisingly, this has led to a severe erosion of belief in the clarity of management's vision of Ames' goals and to a loss of trust in the stability of the line management process itself. Under Syvertson, Ames drifted, rudderless, on the calm seas of benign neglect. These days we find ourselves trying to sail clear of the sharp edged kutlery of that fearful pirate, Cronyism.

The present ADP board is a perfect example of today's response to a management attempt to solve a very real problem. The members, not being blind to this administration's reputation for creating committees at the drop of a hat and then ignoring their findings, have demonstrated very little interest in doing any of the work that is required. And, it being no secret that all decisions, be they by committee, staff, or line management, are subject to being overridden if they give rise to a complaint by a well favored few; people are loath to be involved in making any decision that has implications outside of their own immediate spheres of responsibility. The members of the ADP committee are not incompetent people, but neither are they particularly aggressive or dynamic, and it is from their ranks that the proposed positions will, almost certainly, be filled.

Management at Ames is, by and large, pretty weak above the position of Branch Chief and adding a layer of ADP management at the Division level will not change a weak system into a strong one. The proposed position will almost certainly only be a part time job and will be viewed, by the majority of persons assigned, as just another pain-in-the-neck job to be gotten through with a minimum of trouble.

Ames has never been able to keep aggressive Branch Chiefs from cutting each other up in front of upper management (not even excepting at Washington) in the competition for resources and funds. Is there reason to believe that all of a sudden they are going to become so submissive that they will voluntarily keep in line for someone in the kind of position that this one will become? The only way to give the position any authority is to allow the holder to approve (or disapprove) branch ADP plans, or to allow the person the right to tax the branches for their ADP needs. The last thing that Branch Chiefs need is one more level of arbitrary taxation; and giving weak people the right of approval over important projects has never been demonstrated to work in the past.

The second area of concern involves the extreme interaction of the management problems at Ames and the intense isolation with which the solutions are pursued. One of the most destructive habits that management has developed is the practice of taking the unified result of some study apart, bit by bit, and then extracting out enough dismembered details to ensure that the resulting hash cannot possibly work.

Many, perhaps most, of the areas of concern that have come to be seen by Washington as Ames ADP problems are really attempts to bootleg the solution of some other problem (personnel or procurement, for example) which has become unworkable in its own sphere because the system has failed in a much more general way. Establishing a Division ADP manager will not end this bootlegging of non-ADP solutions. Perhaps not even end the practice of hiding them under the cover of the umbrella of ADP.

I don't mean to imply that we can not reach a decision on Division level control of ADP resources without first cleaning up the mess that is Ames' (or NASA's, or Washington's, take your pick) current management situation. I do, however, believe that it is a non-productive exercise to write a job description for a Division level ADP manager without knowing how more of the pieces of the puzzle will fit together.

How would I catch a fish if this were my can of worms? Clearly Dale and Marcie are the keys to the problem, or rather to the solution of the problem. Dale, because he is the only one in upper management with a secure reputation for strength and integrity, and Marcie because no one else at Ames has demonstrated any vision of how we should manage our computer resources. The difficulty is how to sidetrack Ballhause's meddling with developing plans and how to keep him from giving the whole center away to his buddies. And all this without it appearing that anyone (particularly Dale) is disavowing the Center Director.

The first thing I would do is to disband the ADP committee. I would declare that their assigned task of writing Ames' ADP policy statements had been successfully carried out and thank them publicly and lavishly. This would clear the way of dead wood and might lay the foundation for the restoring of faith in the committee system as a legitimate way to conduct business. The second thing I would do is to form a working group (call it what you like) to design an integrated approach, consistent with those developed policy statements, to all facets of Ames ADP management. I would not make any attempt to have this group be representative of all the research organizations but rather would choose people who are:

1. Willing to work;
2. Capable of taking a center-wide view of ADP;
3. Strong enough to argue management out of cutting their plan to bits, or giving the store away.

Marcie would have to be the head of the group and Code RC would have to provide most of the technical staff work, but they would also need to interface with Communications, RMO, Purchasing, etc. It would be up to Dale to insure the integrity of the process by limiting interference from above

Footnote 3

PRELIMINARY SPECIFICATION OF SCIENTIST'S WORKSTATION

9 July 1987

A. Workstation

1. We see no requirement for a hard disk on the individual workstations but will need a floppy disk drive to allow for the transfer of programs and data.

2. Each workstation must be equipped with a floating point processor.

3. The 4 MB of memory that the SUN workstation which Harry Partridge is evaluating have proven to be insufficient for the kind of work that we routinely do. We will require either 8 MB of memory on the individual workstations or that the file server handle all windowing chores without a perceptible degradation of service.

4. We require the capability to produce presentation quality graphics in black and white. If funding permits, we would be interested in configuring one of the workstations with a high quality color graphics capability for the purposes of evaluating graphical techniques in computational chemistry.

5. The workstations would be required to support the following software or capabilities.
   a. EMACS editor
   b. TEX and a TEX previewer
   c. UNIX operating system (preferably Berkeley)
   d. FORTRAN 77
   e. VT100 emulation
   f. Windowing
   g. Display DIP files
   h. Properly access TELENET and FTP
   6. We foresee the need for a second configuration of the workstation which would be suitable for off-site use. We would require that this system provide the same user interface as the on-site model but also include a 40 to 80 MB hard disk and be capable of driving a small laser printer.

B. File Server

1. We will require a minimum of 200 MB of disk storage on a file server configured to handle 6 nodes.

2. There will need to be some method of backing up the disk storage to an off-line medium.

3. Some thought will have to be given to the physical location of the file server (and potentially several file servers). Space and air conditioning capacity are severely limited in Building 230, but we cannot remove the file server from its workstations to the degree that there is any possibility of reduced reliability of data communications capability.

C. Communications

1. There will need to be 2400 baud modems on the file server, or preferably on each of the workstations, to allow direct access from the off-site workstations.

2. The system needs to be able to avoid accessing either the SYTEK network or the MICOM switches for its routine communication chores. Both of those systems should be available as back up data communications channels.

3. We would support the concept of minimizing the number of interfaces between the workstation and the CCF and NAS supercomputers. If it could be done safely, we would like to see the file servers acting as gateway machines to these facilities.

4. The workstations would be required to access the following facilities.
   a. CCF Supercomputers
   b. NAS
   c. Jupiter VAX
   d. Earth VAX
   e. H. Julian Allen VAX
   f. ARPANET (it would be most convenient if each workstation were a node on the network)

Footnote 4

ASSIGNMENT OF COMPUTER RESOURCES

Other than the mandate from NASA Headquarters there seem to be three justifications for establishing a system for assigning priorities for access to computer resources at Ames. The first, which a minimal degree of monitoring renders null even for temporary and student employees, is that if not watched scientists will play on the computers rather than do their work. Second is the quite reasonable sounding proposition that the scientists who make use of a facility should in some way provide for the upkeep of that facility. And third is to ensure, through the mechanism of accountability to line management, that research deemed important to the nation's welfare is carried out.

Perhaps the best counter that I have heard to the first point is that given by Jim Arnold to a visitor from Headquarters just after the Branch got its first Terak personal computers. The man (fresh to Government Service from the world of Academe) warned Jim to watch his scientists carefully or they would spend all of their time playing "Alien Invasion". Jim replied that, while this might be the case with college students, his experience was that professional scientists are what they are because to them research is vastly more satisfying than any computer game.

The second point seems to lead naturally to the concept of collecting Code RC's operating funds through the establishment of some kind of profit center. I believe that it would be a mistake for Ames to return to this mechanism. The problem, as we have seen, is that the method fails disastrously if the Division does not, for whatever reason, cover all of its costs. Therefore the financial managers for Code RC will once again be forced to spend a great deal of effort to see that their accounting algorithm is conservative (that they will always make a profit) and that it is collecting money at the rate that they require. None of this will make sense to the user community who will take the oxymoronic view that the Computer Center is making an unjustified profit at their expense and that it is all just pretend money anyway. Code RC management should, as now, be guaranteed the funds they need to run their Division regardless of any assignment of computer resources to the research community. Code RC staff must, however, be an integral part of the prioritizing system because it is they who must design the metrics that will be used to measure usage and the units that will be rationed out.

As far as the third point is concerned, upper level science line management at the Center has often been somewhat loath to make direct judgments about the quality of research projects. One of the factors in this stance is that the breadth of technical expertise that a manager brings to the job necessarily becomes too narrow to cover the work he must oversee as he moves up the management ladder. The issue of accountability is generally felt to be satisfied by the advocacy process that results in the disbursement of supporting funds from NASA Headquarters. The fact that these kinds of value judgments can be successfully made at the local level is demonstrated by the workings of the committee that makes the recommendations for distribution of the Directors Discretionary Fund. Promotions in grade are another area in which individual research projects are rated as to value, and the multi-level promotion board process, while perhaps overly comprehensive, might serve as a model for a method of assessing potential resource priorities. Just as it is the individual Branch Chief who sets priorities for his scientists by lobbying for their financial support, so too his is the basic responsibility for assigning priorities for resources available to his researchers and then for advocating their positions to whomever ultimately makes the assignment. It also seems clear that this ultimate responsibility, at least in the area of computer resources, lies with the Director of Code R. His organization controls both the computers and their heaviest users. Obviously there must not appear to be any prejudice for Code R projects at the expense of those of the other research directorates.

One of the lessons we should have learned from the failure of the operation of the CCF as a profit center is the danger of losing the support of the user community. The individual scientists and their managers

must understand the basis and operation of the priority assignments well enough to believe that they, and everyone else, are being given a fair opportunity to do their work. Secondly the users must never get the idea that computer resources are being wasted because of the rigid enforcement of the priority system. It is in this area that any proposed priority system will face its severest test. Computer resources are a very tight function of time so that a CPU second unused is one lost forever. Some of the users may well not be able to make use of their share of the resource evenly throughout the year, which will cause scientists who have been given a limited budget of resources to play with strategies to beat the system. They might use up their budget early in the year in the hope that later machine usage will be light. Organizations will bargain away unused resources thus defeating any rationale behind the original assignments. Scientists who are facing a limit on their usage will make arbitrary, and unqualified judgments about the value of the work done by other people. Students are already a constant and continuing target of criticism by researchers, criticism that may or may not be valid but is certainly done with the a priori assumption that work by students is of little value. In another area, we all remember the criticism that the assignment of priorities for NAS usage got when it was discovered that someone had used his Cray time to generate Pi. It is conceivable that such a project could be quite legitimate, perhaps part of a test of the machine's functioning. But in a situation where people compete for resources the assumption will be made that this computer time was wasted, accusation will take the place of reason and the system will be continuously challenged. It is inevitable, in any situation where there is competition for a resource, that people will challenge each other's allotments. If the allocation scheme is not seen to be orderly, firm and fair (as is the case now with project funding and promotions) then ill feeling and maneuvering to get around the system will nullify any management value the scheme has. Jobs that are run on the science computers at Ames might be divided into two broad types, each of which may well benefit from having a separate set of metrics. One type of job resembles a product that might be sold to someone else. These consist of stable codes (probably the result of a successful research project) that are run in a production mode and produce the answer to a query from another NASA center, another agency, or any outside source which might be expected to transfer money to Ames for the services rendered. The other kind of job consists in code development to further some research project. These jobs bring funds into Ames because someone, usually the Branch or Division Chief, has convinced a money source that the work is worthwhile and that his researchers can successfully do it. If management does not choose to discriminate between these types of jobs in choosing the priority algorithm at least some effort should be made to see that neither type is heavily discriminated against.

The following are a few obvious discriminating factors that might be used to assign priorities to research projects. Resources might be distributed as a function of:

1. The number of dollars that a project brings into Ames;
2. The number of tax dollars that a project or Branch pays to the Ames Administration;
3. The number and quality of scientific papers produced by a project;
4. Any time limit that pertains to the project such as scientific meetings, interfacing with scheduled events (as perhaps space shots) and deadlines established by management
5. Past usage, although this presents some problems. An uncritical acceptance of past usage as a guide to future allotments may well lead to the automatic continuance of projects that are in reality only constant in their failure to produce results. The urging that "What's past is prologue" was spoken by one scoundrel to another in an effort to incite to murder. I will speak about the value of examining the history of a group's computer usage again later.

In practice these discriminating factors might be used separately, in any combination, or be replaced with something totally different but in any case the mechanism of assigning resources and the rationale behind it must reflect in some visible way the vital interest of each of the diverse groups of computer users at the center if Ames is to enjoy any peace at all.

The next, and perhaps the most sticky, subject is how the system will be enforced. It would be the easiest course to make Code RC responsible for policing and enforcing the priority system. And it is clear that they will have to be responsible for reporting usage in some meaningful way to whomever does have to

enforce the priorities. In my opinion it would be a very short sighted policy to require Code RC to shoulder this responsibility for the following reasons, any one of which I deem to be sufficient.

1. It is enough of a job to run the Division as it is without having to bring in extra staff to monitor and enforce what, in this situation, will become a real nest of worms.

2. Enforcement would reestablish a friction point between Code RC and the user community that we have all worked to remove.

3. It is unfair to put Code RC in the position of passing judgment on the value of science projects. To do so gives any scientist who wants to protest the size of his ration a weapon that can't be countered and so the Division will waste an immense amount of time and energy defending itself.

Computer resources, having been divided up by some method, should be assigned to the Branch Chiefs rather than to each individual research project. And it should be the responsibility of the Branch Chief to see to it that the researchers within his organization stay within the limits imposed by management. Virtually all basic management goes on at the branch level and all of the resources available to a branch are the responsibility of the Branch Chief. My feeling is that no Branch Chief should be called upon to justify the use of resources within his organization once those resources have been assigned. If, in a Branch Chief's judgment, all of the branch's computer time should go to students then that judgment should not be open to question by the user community at large. The ultimate responsibility for the successful completion of a research project lies with the Branch Chief, and it is on this level that his judgment must be proven or defended. It would only be appropriate to censure or question a Branch Chief if he consistently or blatantly overran his allotment. To limit a Branch Chief's flexibility in marshaling and dispersing the resources at his disposal will both eliminate serendipitous research and hinder structured research.

No matter what mechanism is set up for assigning resources I would most strongly recommend that there be enough latitude left in the system to allow for some discretionary use of the computers. It is clear that this addition to the policy will complicate the task of fairly distributing these critical resources, however, my feeling is that assigning 100% of CPU time to existing or expected projects at some arbitrary time during the year would be placing a great handicap on productivity. Some of the requirements that I can foresee needing discretionary time are:

1. Code RC staff will require time to engage in their own examination of matters relating to the use and operation of computers at Ames. If Ames management chooses not to accept the idea of discretionary time then Code RC should be provided with computer time to use as its staff thinks best;
2. Computer time might be a useful addition to the resources available to the Directors Discretionary Fund;
3. If responsibility for computer usage is to be assigned at the Branch level there are people, like Division Staff Scientists, who will drop through the cracks without some further provision being made for them;
4. Not every estimate of requirements will be accurate nor will every assignment of priorities completely cover the actual need. Therefore, there will be a requirement for some mechanism to adjust resources after the initial assignments have been made.

It should be noted in conclusion that there is a de facto priority system in operation at Ames today. It is based on the energy and enthusiasm of individual scientists. Those researchers who are willing to submit jobs to the computers from home, during off duty hours, on weekends and on holidays will get more computer time. Scientists who are willing to do the extra work to get codes ready in advance of the arrival of a new computer will get more computer time. Those willing to structure their codes to fit an unpopular computer will get more computer time. Because the research environment at Ames is so dynamic, past usage by an organization is not a good metric to use by itself. However a careful examination into how an organization has used the available computers in the past can be very revealing as to the spirit of the scientists and to the quality of their management. If the assignment scheme

ultimately decided upon severely limits scientists of this caliber they will first try to circumvent the system and if unsuccessful will leave Government Service. No system of resource assignment is worth enough to be bought at the price of keeping our best people.

Footnote 5

Use and Funding of the Central Computer Facility Cray

The Aerophysics Directorate provides virtually all of the funds that are required to support the Central Computer Facility Cray Y-MP8/832.

Total Funding of the CCF Cray

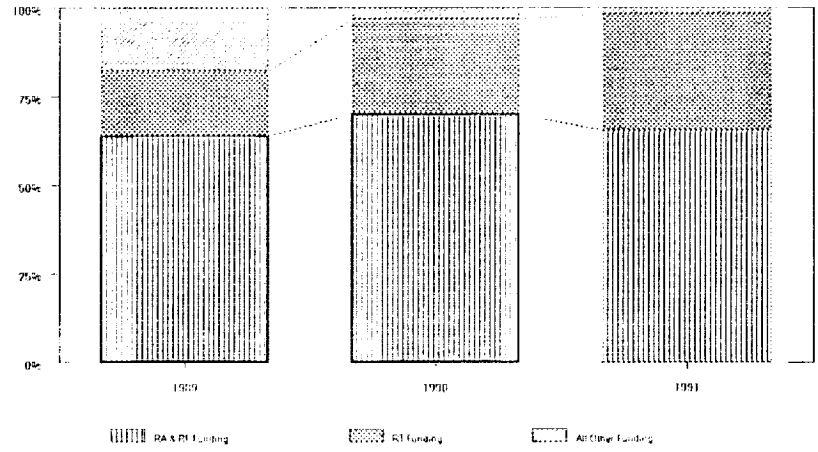

NASA Headquarters' support of the research of the Thermosciences Division, through RTOP funds, is strong and is increasing. The three RTOPs that the Division uses to pay for its computing resources are Materials - 506-43, Aerothermodynamics - 506-40, and High Energy Aerobraking - 506-42.

The amount of money that the Division is required to pay for its use of the CCF Cray will have been more than doubled between 1989 and 1991.

Code RT Funding of the CCF Cray

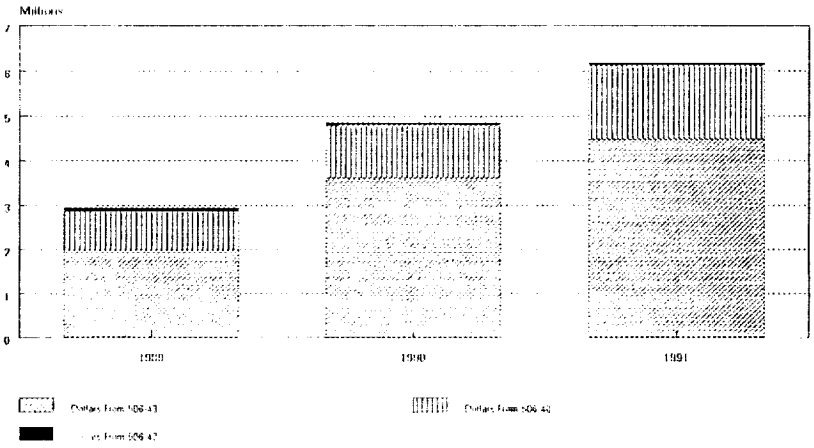

Just as the total number of dollars spent by Code RT is increasing, so too is the percent of the Directorate's funding that the Division carries.

Code R Funding of the CCF Cray

by Percent

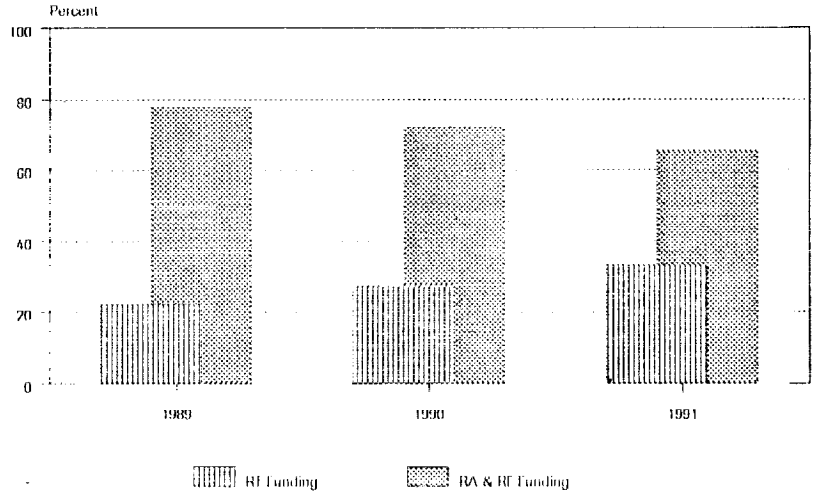

While the Division's share of the costs increase by a factor of 1.5 its share of the resource decreases by that same factor.

Code R Use of the CCF Cray

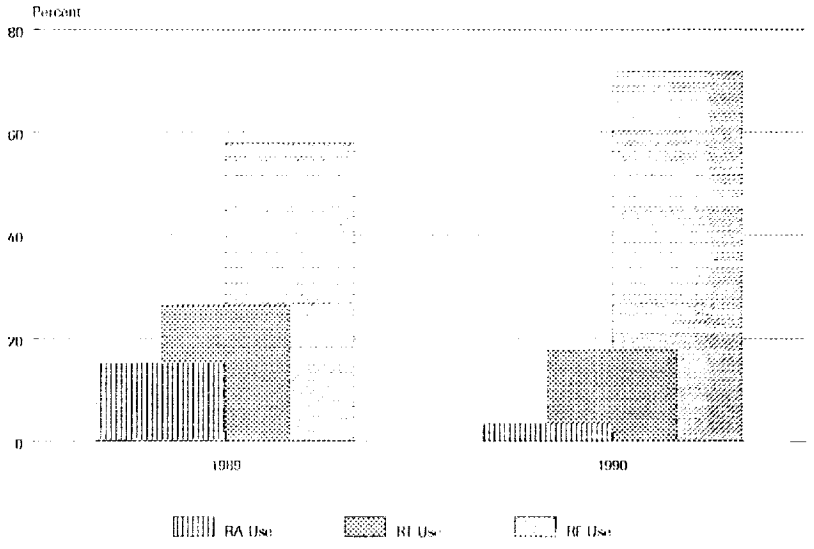

From 1985 until 1988 we were able to ameliorate the shortage of time on the Cray by having access to (and paying for) a significant part of the time available on the CCF CDC 205.

Code RT Use of the CCF CDC 205

All CPU Times are Converted to
Cray XMP/48 Time and are Expressed
As a Percent of Total Cray Use

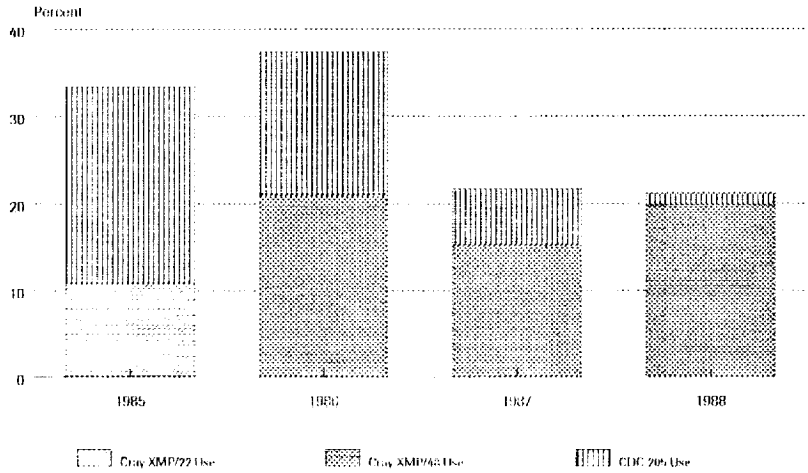

The only other major computational resource available to the Division is, of course, the NAS. Last year, with 20 Research Project Groups, the Division's scientists used 53.6% of their allotted time on the NAS. This year the Division's 12 Research Project Groups have used nearly half of their allotment with 31 weeks left in the computing year. Each of these Research Projects represents a major effort that could not be attempted without the enormous resources of the NAS. None of these projects is sufficiently broadly defined nor sufficiently richly endowed with CPU time to enable the Division to use the NAS in lieu of the CCF Cray. For good or ill the CCF Cray is the Thermosciences Division's bread and butter.

Code RT Use of the NAS 1989-1990

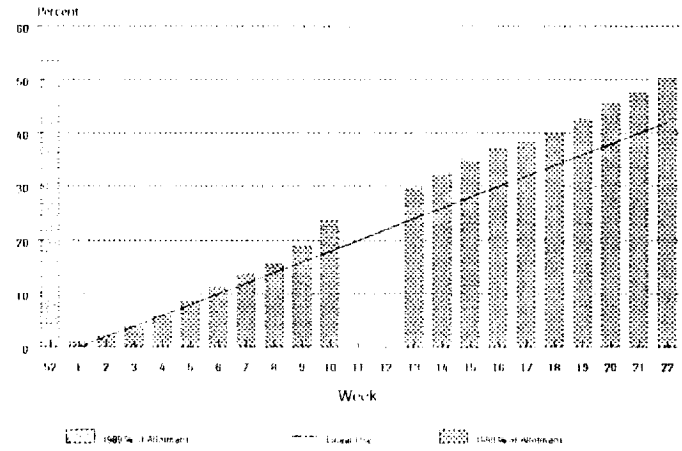

The heart of the problem lies in the number of user accounts that have been assigned. There are 756 accounts on the CCF Cray YMP, as contrast, the NAS, which is a national rather than a local facility and which spreads its work load out over 2 supercomputers, has 555 accounts. The relationship between the number of accounts and CPU use on the CCF Cray was brought to the Director's attention earlier this month by Code RC. The following plot is presented with their kind permission.

Y/MP USE BY ORG CODE

FY 90

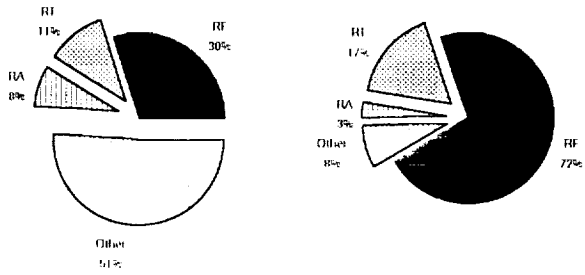

NUMBER OF ACCOUNTS CPU USE

Plot Courtesy of Code RC

The reason that the imbalance in the number of user accounts provides such an overwhelming advantage is that no user account may have more than 3 jobs in the queue at one time. Just how big the advantage is, is attested to by the disparity in the number of jobs that each organization has been able to run. Code RF runs more jobs than Code RT by a factor of 5 and more than Code RA by a factor of 14.5.

Code R Use of the CCF Cray

by Number of Jobs Run

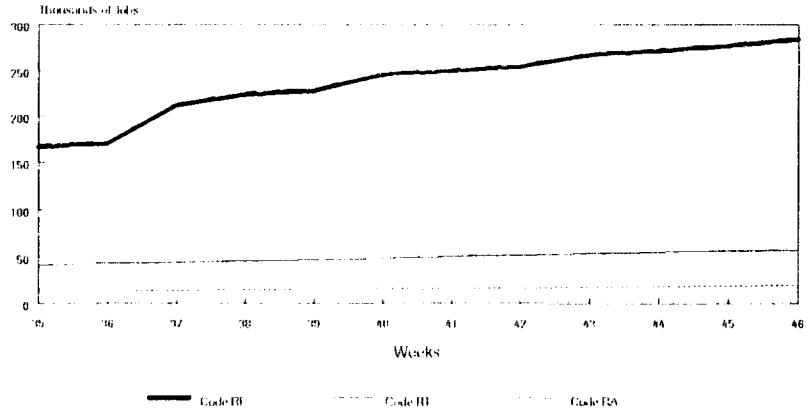

We need to relate the problem to the real world. I was mistaken in thinking that we could use missed deadlines to make our point. The problem with the past is that when we tell Ron about our failures he will ask us "Did you talk to Marcie about your problem?" or "Did you bring your problem to higher management?" There are no satisfactory answers to either of those two questions. Either Vic or Marcie, or most likely we, will end up looking bad, and none of those things is good for us. There is another problem with rehearsing our insufficiencies in public. Management is not given the job of missing

deadlines but rather is to see to it that deadlines are met. You never have to explain your successes, for no one questions success. You never have to explain your failures, for such explanations are invariably viewed as excuses and it is your failure rather than your excuse that is remembered.

There are other traps to avoid in any exercise that strives to correct some imbalance in the system. We must not give the impression of being whiners who have come to complain about all of the bad, bad people out there who are doing us wrong. The Turbulence people thought, just as we now think, that they did not have enough access to the Cray. They brought their problem to the attention of higher management and were given a privileged queue. That is just what we are doing. DON'T ADDRESS TO PERSONALITIES. DON'T ATTACK A MANAGEMENT SOLUTION THAT WE MAY WANT TO USE OURSELVES.

And it is most important not to alienate our friends. WE HAVE NOT COME TO COMPLAIN ABOUT THE SERVICE GIVEN BY CODE RC. Marcie can be a strong and valuable ally if we give her the opportunity.

We can't base our argument totally on a question of money, although my view of the importance of this issue is reflected in the high number of plots relating to funding in this presentation. The problem with money as the overriding issue is that not all of Code RT's branches are equally well endowed. I can't believe that an organization that is well funded because of the labors of its management and because it possesses a history of success is going to be willing to fund computer time for the entire Division. Nor do I believe that the NASA Headquarters sources of its funds will long permit it to do so.

We probably don't want to allow the discussion to dwell on events and personalities at the branch level. If we can keep the focus on the Division we may be able to cast a protecting wing over any weaknesses that we may have. It would be a mistake to forget, however, that the Ames' Community is a very small one and that no one fools anyone (particularly a Director) for very long.

There are possible justifications for asking for a greater slice of the supercomputer pie other than money. One would be to express concern about our capability to meet some specific, important deadline in the near future. This deadline should be one that Ron cares about and one that we are certain that we will meet. It would be very embarrassing to get everything we ask for and then fail to deliver, and neither all of our projects nor all of our people are equal in their potential for success.

Another line of reasoning, and one that strikes directly at the problem of Civil Service vs. students, is to point out that, given what the Government has to offer as an employer, the only way we can attract and keep scientists of any potential is to allow them to attack the most interesting problems available and to give them access to the best possible tools for the solution of those problems.

This is where we need to tell Ron what we want him to do. Our recommendations need to be specific and coordinated with Code RC.

Be careful not to try to set up Code RC as a policeman. This must be an easy mistake to make because we seem to keep making it.

The following possible solutions have been offered, I report rather than advocate them:

1. Restrict off-site users to a fewer number of jobs in the queue at one time;
2. Tie non civil service off-site users to an on-site civil servant and make them share some quota;
3. Reduce the number of user accounts;
4. Request a separate queue for Code RT users.

Footnote 6

Allottment of NAS Resources to Code RT

1989-90 Computing Year

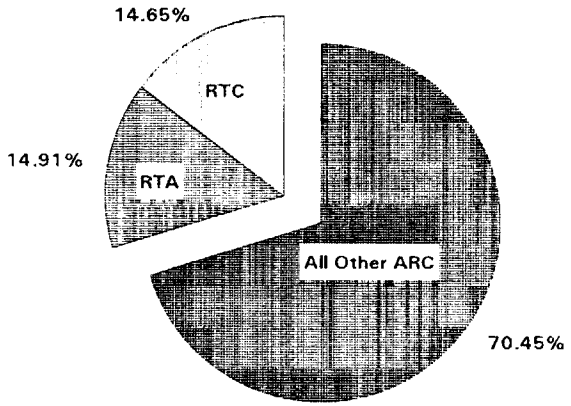

1990-91 Computing Year

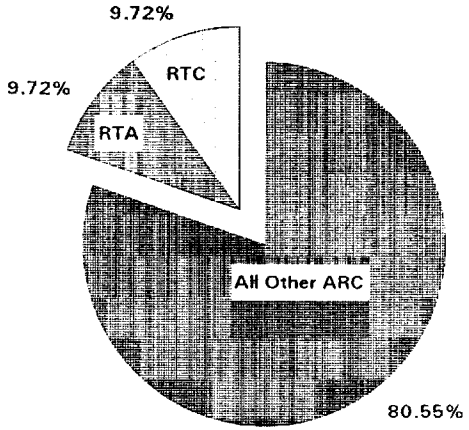

1991-92 Computing Year

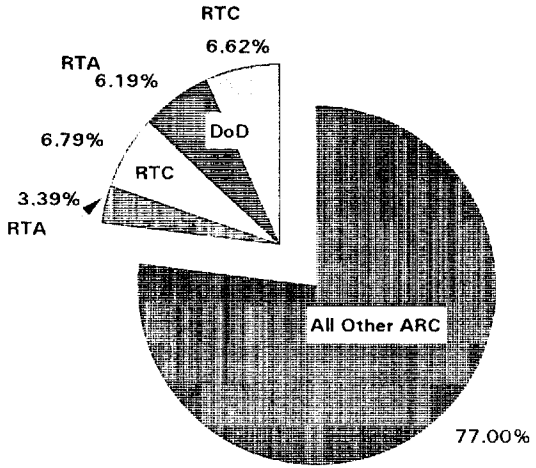

Total Number of SBUs Allotted to Code RT

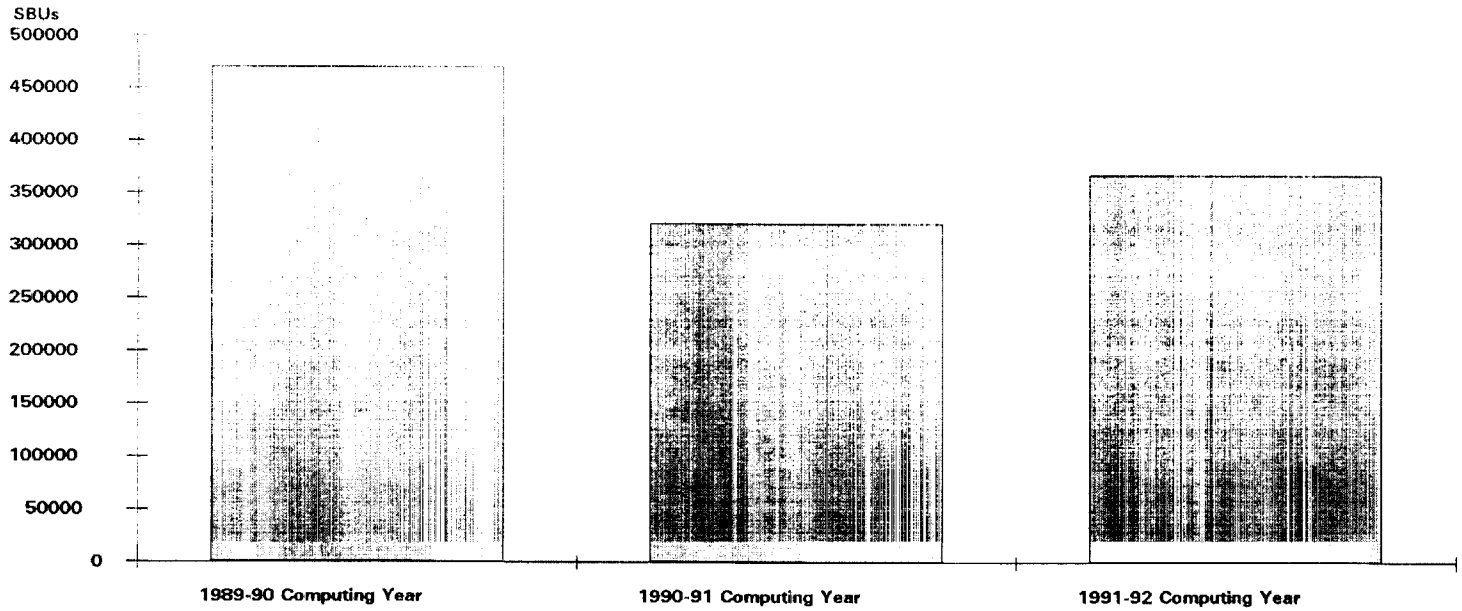

Code RT SBU Assignment as a Percent of the Total Ames Assignment

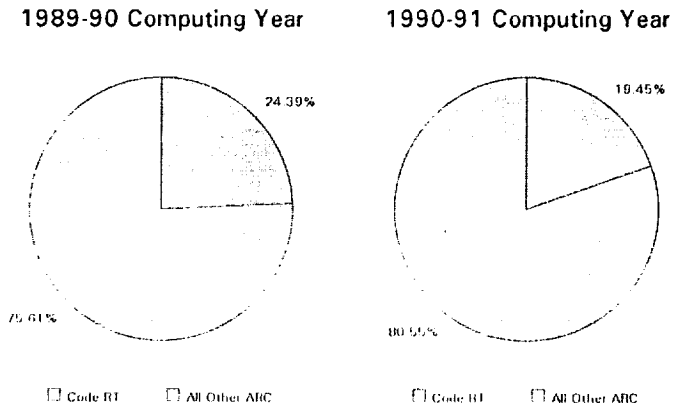

Total Number of SBUs Assigned to Code RT by Computing Year

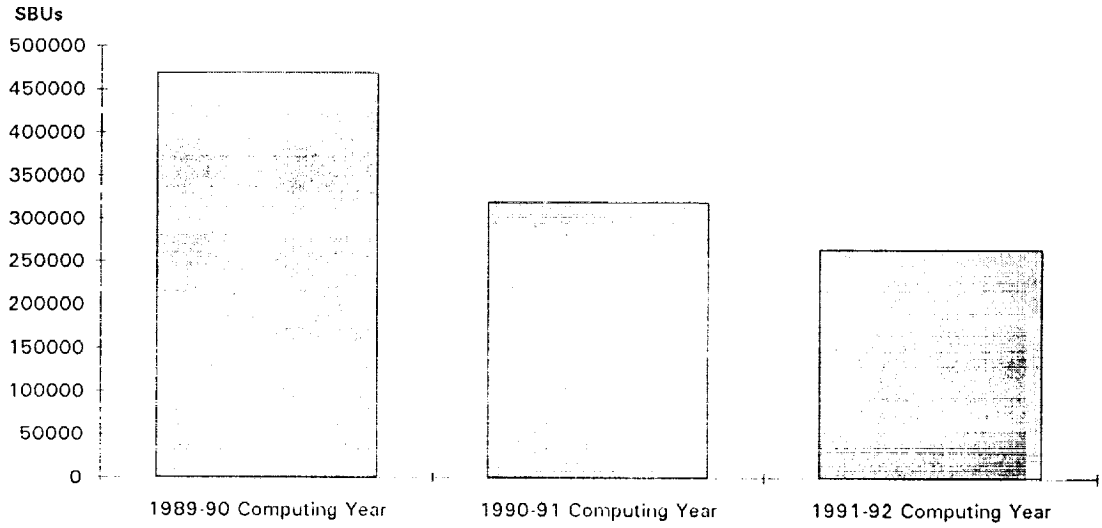

Code RTA SBU Assignment as a Percent of the Total Ames Assignment

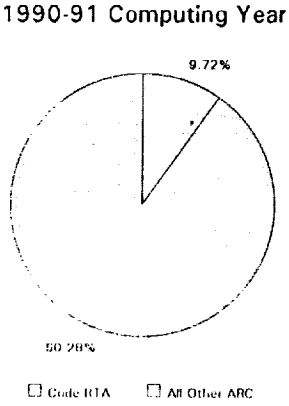

Total Number of SBUs Assigned to Code RT by Computing Year

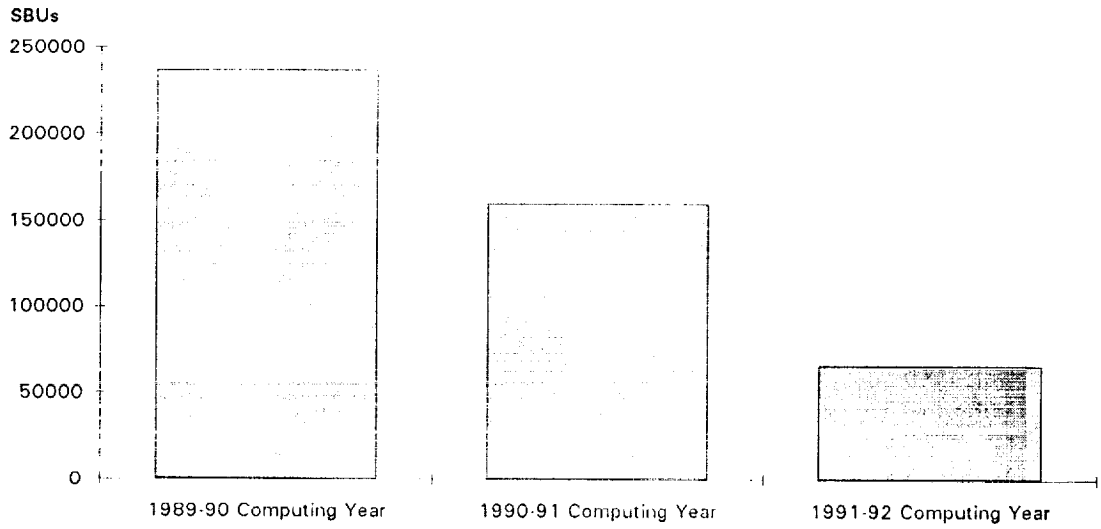

Code RTC SBU Assignment as a Percent of the Total Ames Assignment

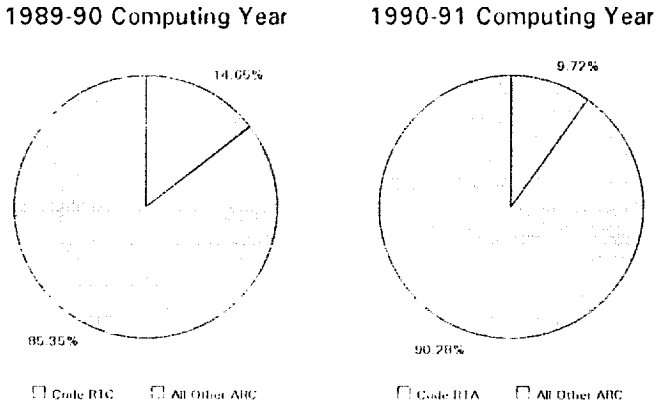

Total Number of SBUs Assigned to Code RTC by Computing Year

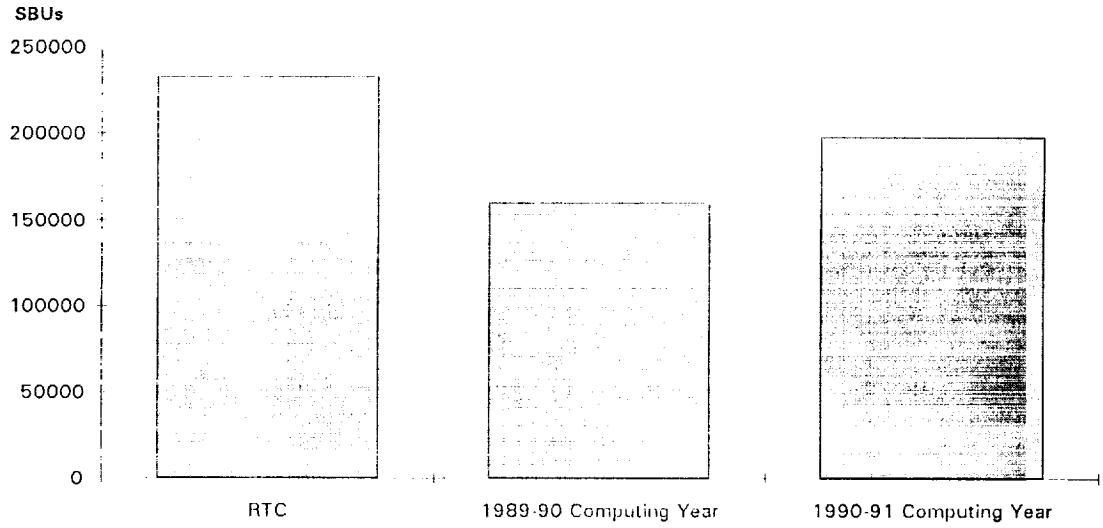

Code RTA NAS Project Continued from 1990-91 Computing Year
NAS Project 2010 - AFE Flowfield Simulation - Feiereisen

NAS Resources Requested and Assigned

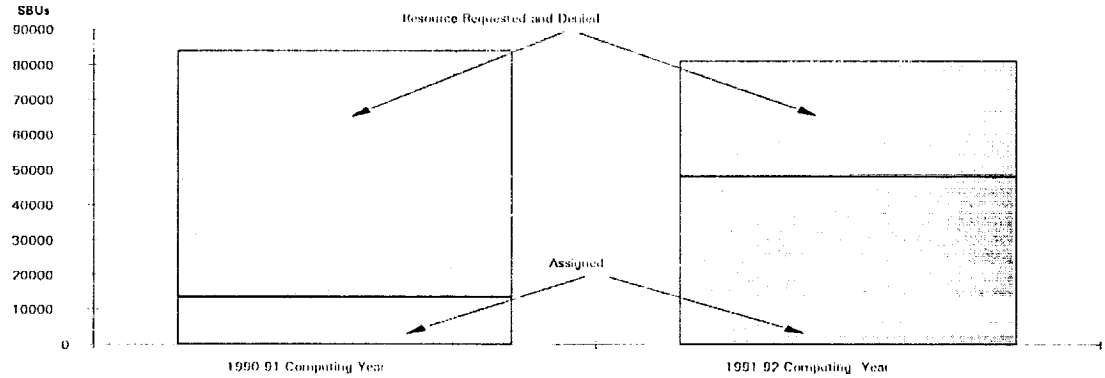

Actual Use vs. Linear Fit to Total Assignment in Each Year

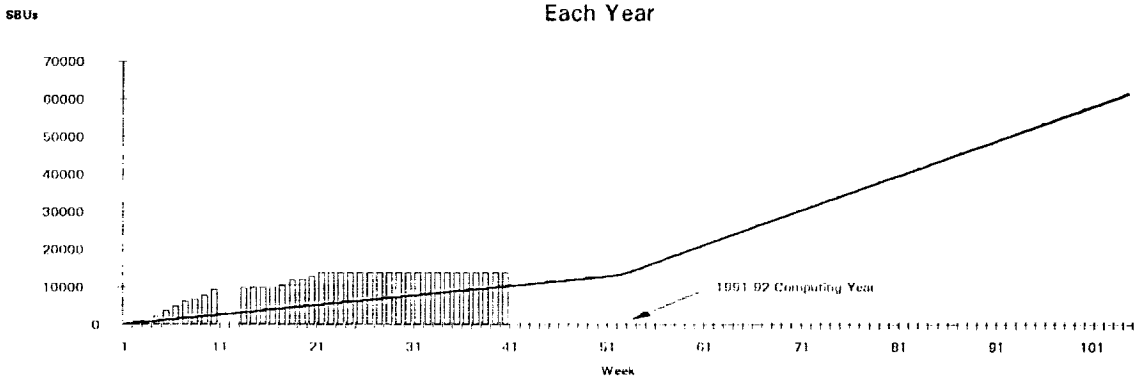

Actual Use vs. Linear Fit to Total Assignment Over Both Years

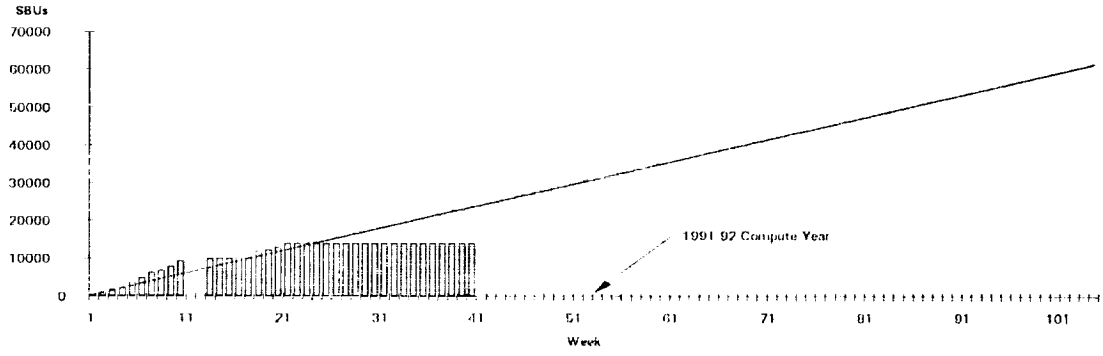

Code RTC NAS Projects Continued from 1990-91 Computing Year
NAS Project 2019 - Chemical and Physical Properties of Propane-Air Mixtures for the High Speed Research Project -Jaffe

NAS Resources Requested and Assigned

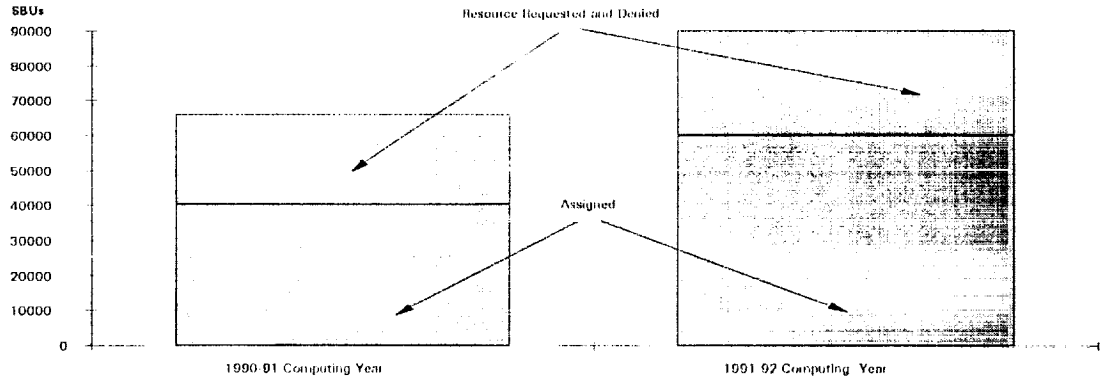

Actual Use vs. Linear Fit to Total Assignment in Each Year

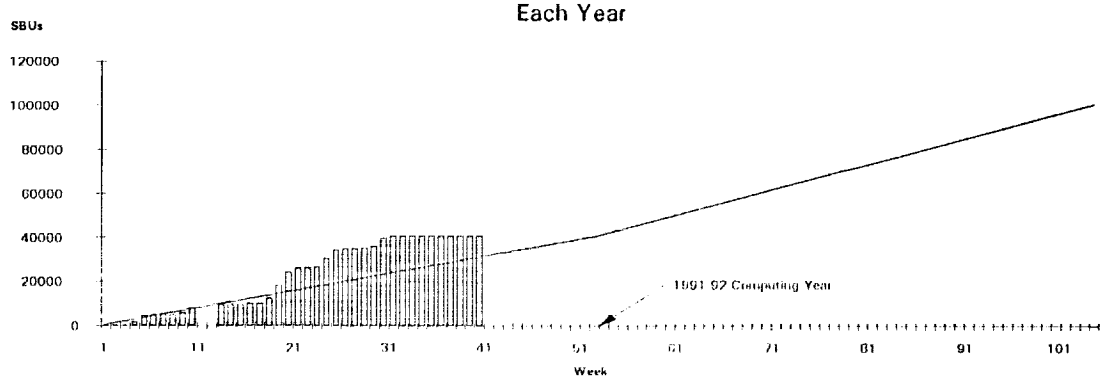

Actual Use vs. Linear Fit to Total Assignment Over Both Years

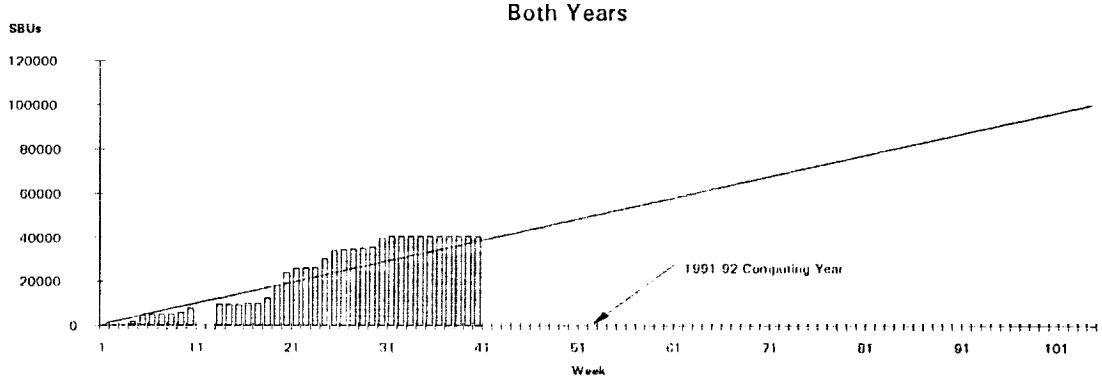

NAS Project 2036
Molecule - Molecule Interactions Important for Aerothermodynamics Flowfield Calculations - Partridge

NAS Resources Requested and Assigned

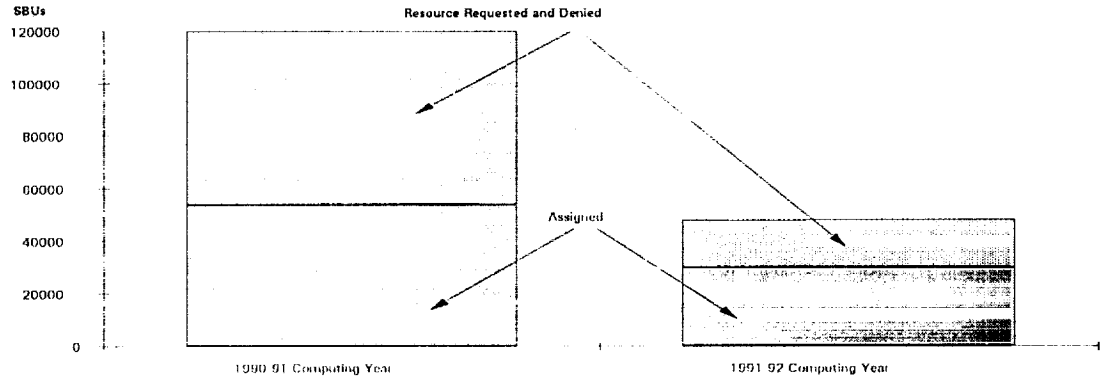

Actual Use vs. Linear Fit to Total Assignment in Each Year

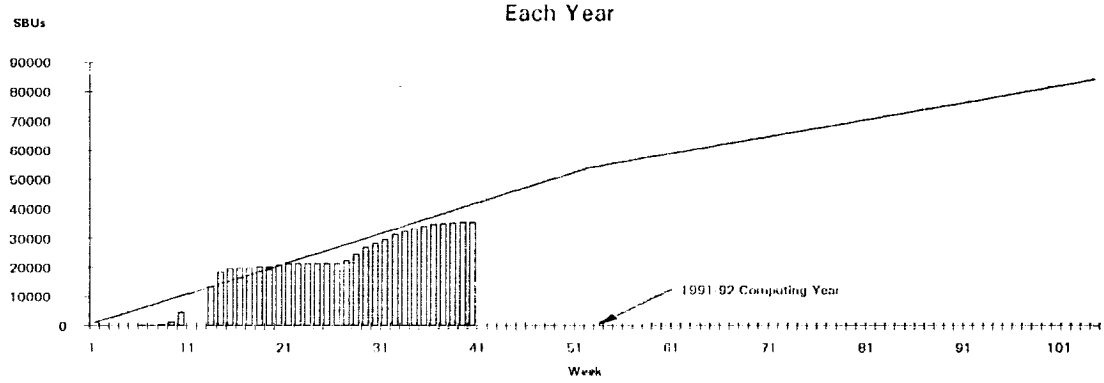

Actual Use vs. Linear Fit to Total Assignment Over Both Years

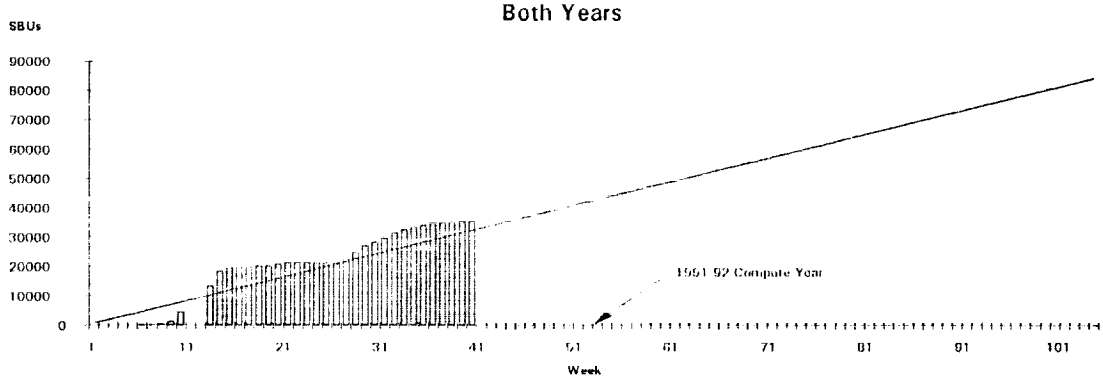

1989 NAS Project 1234
Hydrogen-Air Chemistry for Hypersonic Vehicles - Schwenke

NAS Resource Requested and Assigned

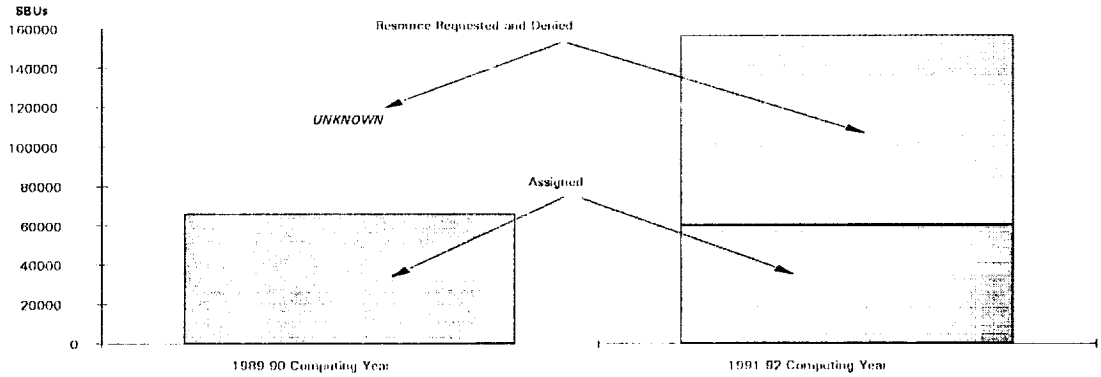

Actual Use vs. Linear Fit to Total Assignment in Each Year

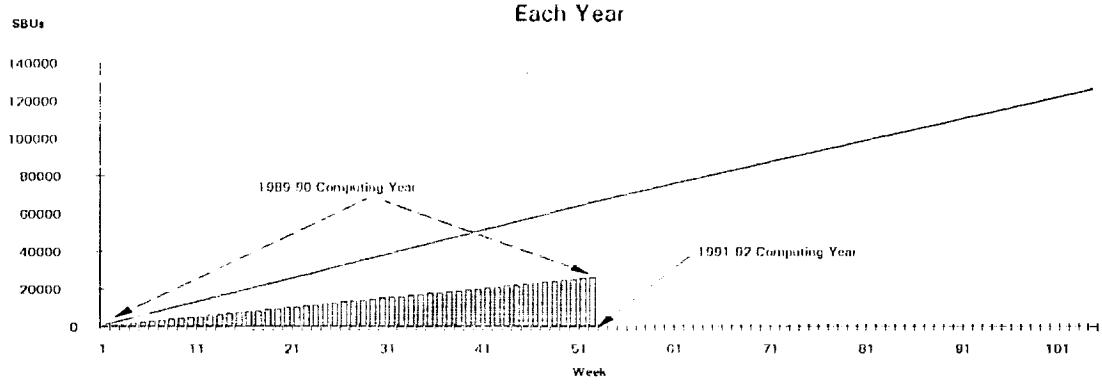

Actual Use vs. Linear Fit to Total Assignment Over Both Years

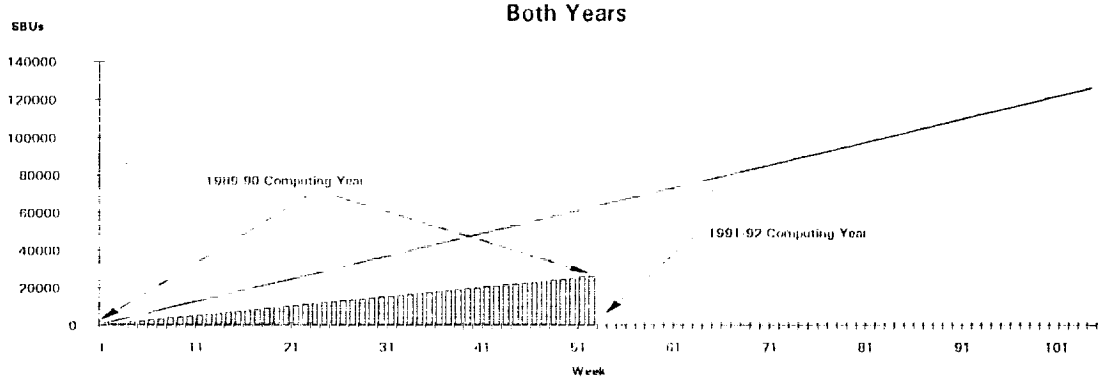

Code RT NAS Projects That Are New this Computing Year

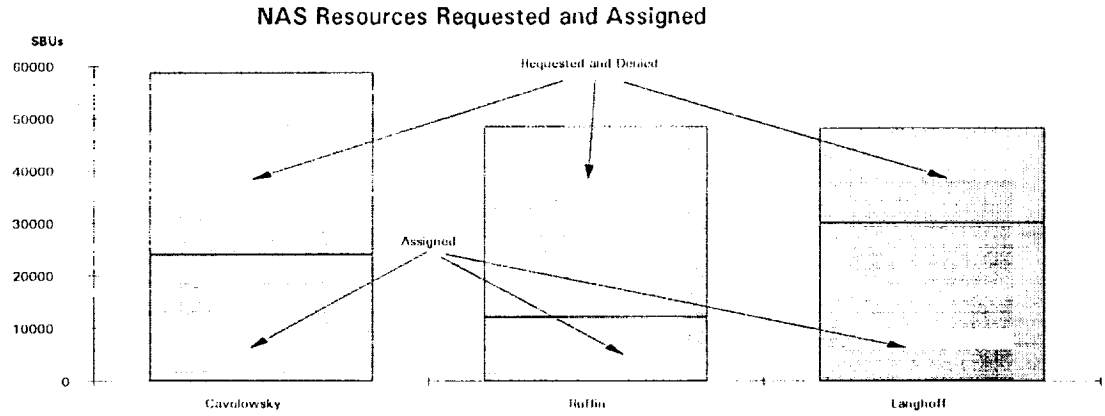

The Effect of Adding the 400 Discretionary CPU Hours to Each Project

Code RT Projects Accepted by the NAS

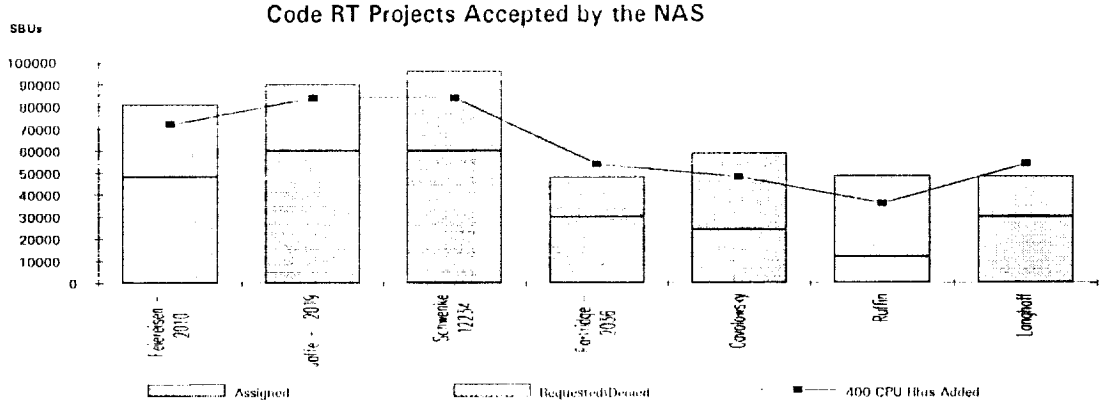

1990-91 Projects That Were Not Continued

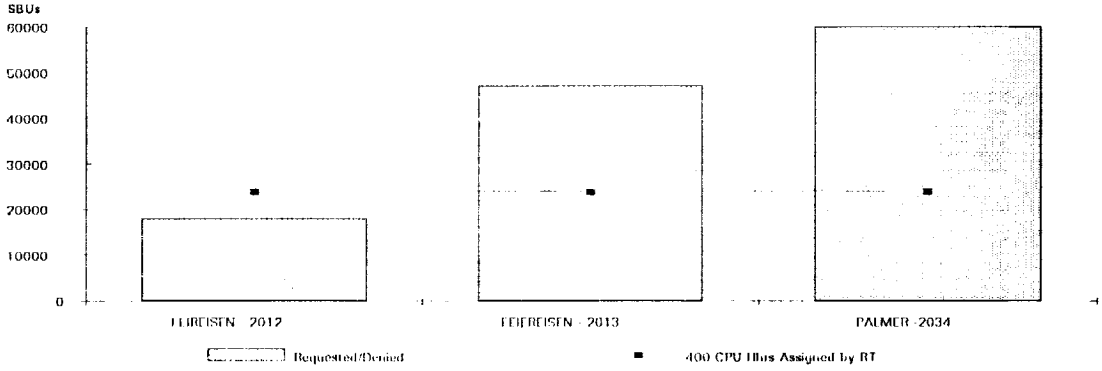

Newly Proposed Projects that Were Not Accepted

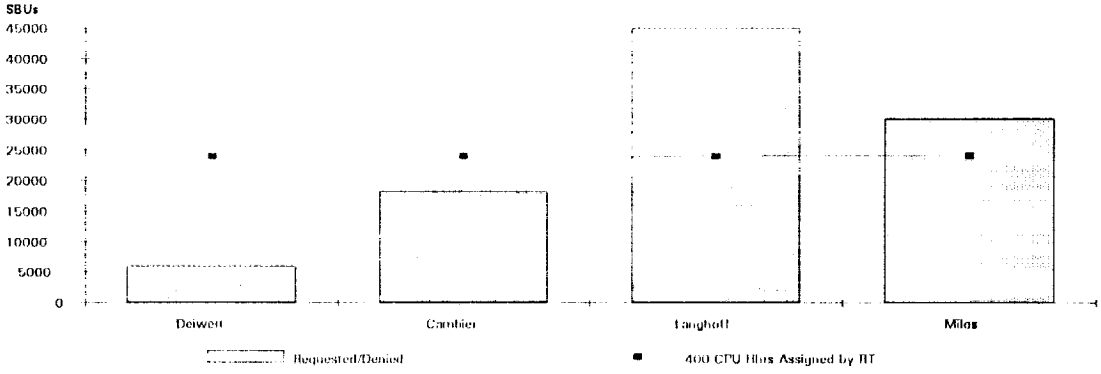

Footnote 7

During the reporting period. Sophie Duckett worked and researched the following project : Information Distributions Systems. As a result of her work a couple of reports were written which follow

Information Services Committee Report

May 21, 1992

Introduction

A great need exists for exchanging information among customers of the Computer Systems and Research Division. The Division recognizes the importance of providing timely and accurate information, and understands that this exchange can greatly affect the productivity of Ames researchers. Recent informal analyses, such as those undertaken by the User Needs Assessment team, confirm that users rely on us not only for information about systems supported by the Division, but also want and need an uninterrupted flow of reliable, consistent technical information covering much broader areas.

The decision to reevaluate the methods of disseminating information was prompted by several events, including the cancellation of a major instrument of communication, the *on_line* newsletter, and the realization that the Computer Information Center (CIC) no longer fulfills important user needs. These conditions, along with the recognition that information dissemination is important to the user community, have prompted the Division to look for new ways to deliver time-critical and important information.

1.0 Purpose of this Report

An Information Services Committee was formed to take the following actions and report findings to Division management:

- Review existing processes of information dissemination
- Define the types of information required by the user community
- Identify target communities to disseminate information
- Recommend methods of dissemination

Each of these tasks are discussed in this report.

2.0 Existing Processes of Information Dissemination

Current methods of distributing information are listed below, with a brief analysis of the advantages and shortcomings of each. Most of these methods are generally available to our customers during regular business hours.

- The telephone is the most frequently used and vital way to work though technical problems with users. User Services also contacts resource monitors and branch chiefs to relay time-critical information. It is often the only means available for distributing information quickly to specific members of the user community. Users can leave a message at any time, however the phone is only staffed 8:00 A.M. to 5:00 P.M., Monday through Friday.

- *News* announcements are provided on various computers, allowing users to read news and print output. Announcements cover topics such as facility and computer maintenance and/or upgrade schedules. News announcements are only available if users log on and invoke *news*. When computers are down, there is no way to communicate that fact, or to report computer status changes.

- Electronic mail is useful for relaying information after regular business hours or to communicate with people who are hard to reach. It is also an invaluable means for communicating

difficult problems to User Services. Users can e-mail faulty programs, along with the output, to User Services for analysis. E-mail is a well understood, informal, and effective way to distribute information between individuals or from one person to a group of people with common interests.

- User guides developed by User Services vary from one to three volumes of over 400 pages each, with a distribution of 400 to 1000. They are distributed through the Ames mail system to a broad spectrum of users whose abilities span the range from novice to expert user. These guides are generally well received by the user community. In this rapidly changing environment, user guides require frequent updating, which is currently done annually (or less frequently) due to the labor and printing cost associated with distributing these large volumes.

  Some smaller documents developed outside of User Services do exist on line. Man pages, a form of documentation, exist on line, as well.

- Technical reports and publications are articles submitted for publication internally to NASA and externally to professional societies and industry consortia. Although access to these reports is not restricted, they are usually only distributed to a few interested individuals within the Division and are not made available to a wider audience.

- Hardcopy mailings are used to notify users of hardware and software updates and training classes. Mailings can be targeted at specific groups, or mass mailings are done to notify all users of important changes. While these mailings are adequate for timed and predictable changes, emergency or other time-critical changes cannot be handled this way, as a minimum one-week lead time is required for the information to be adequately distributed.

- Meetings with the user community play a key role in communication among diverse groups. They provide a method to efficiently disseminate information to large audiences and

have the additional quality of personal contact with the customer community. Annual vendors' meetings inform Division staff and interested users about the future direction of technologies.

- The User Advisory Group, consisting of representatives from the user community and the Division, has become a valuable forum for information sharing. The charter of this group is to discuss the needs and critical requirements of the users. The group meets on a monthly basis and recently began distributing meeting minutes.

- Training classes are conducted by User Services and Lurnix. Classes are selected based on input from previous class evaluations and user surveys. These classes are self-paced lab-lecture or lecture-only format. Workstations are available in User Services to continue self-paced training. Students who are unable to attend classes can receive training material to practice on their own, or can get assistance from User Services.

- Workshops and presentations on specific topics are presented on an as-needed basis.

- The Computer Information Center (CIC) is a repository for manuals, periodicals, books, and Ames publications, to which users may come to research technical information. The CIC grew out of the need to order and store manuals from computer vendors. In the past, our customers found this a valuable service; however, the need to order manuals has dropped off dramatically due to the ability to purchase documentation through the IAS contract and Ames' permission to duplicate Cray manuals. Many of the CIC functions, outside of ordering manuals, duplicate those of the main library. The CIC does provide several clerical services that are not performed elsewhere, such as maintenance of mailing lists and routing of publications and periodicals. The CIC itself, since it is not systematically maintained, has almost no users.

3.0 Types of Information Required by the User Community

The types of information that will be most useful to our user community have been identified through discussions with users over the last year. This information was obtained from the User Needs Assessment activity, the User Advisory Group, and CCF User Services. These groups have yielded a significant amount of requirements information. The following three sections identify some of the requirements expressed within these groups.

3.1 User Needs Assessment Team

A recurring theme that the User Needs Assessment team discovered (beyond the well documented "more, bigger, faster" supercomputer requirement) was the request for more information to help them effectively use our systems and services. Without exception, every user organization interviewed so far (RAA, RAC, RAF, RFR, RTA, and SL) suggested that we can help improve user productivity by providing them with such information. Currently, we satisfy information requests and information dissemination, with varying degrees of success, using the methods described in Section 2.

The user community not only wants the Division to be "information brokers," they want us to take an active role in providing information that will be useful in developing new technologies, or developing applications that will utilize emerging technologies. A list of the types of information which may be of interest to the users is provided in Attachment 1. More specifically, users want us to provide information on the following subjects: code optimization methodologies and techniques; massively parallel systems and the applications methodologies to take advantage of these systems; image processing; graphics; visualization; and "real-time" Unix. In addition, users would like us to sponsor conferences on, for example, code validation methods and techniques, image processing, and software selection methodologies, which would expose them to new techniques, and consequently increase their productivity.

In addition to these types of information, the user community suggests that, at a minimum, we provide information in a standard fashion that is easily accessible, reliable, and well-maintained. They want to be able to see, order, or print this information at their discretion. They do not want to rely on getting their names on multiple lists or subscribing to internally developed publications. They don't want to leave their offices to search in multiple places for information about the CCF, and prefer to have one standard way of accessing information, not two or three on the multiple systems where information resides (NAS, ACF, various file servers).

The user community is also looking to the Division to provide standards for such items as:

- information transfer
- formats for documentation transfer (specifically, MacIntosh-Unix sharing of information)
- user-friendly interfaces to the information
- features to access and move the information (such as searching, printing, forwarding)

3.2 User Advisory Group (UAG)

One of the recurring issues in the UAG meetings is the availability and reliability of certain types of information. UAG customer representatives have indicated that they rely on computing-related information provided by the Division to successfully conduct their research.

Users were concerned when *on_line* was cancelled, because they relied on the technical information that the newsletter presented. They have requested a replacement mechanism for distribution of that information.

Users would like us to respond quickly to their information requirements. They want to see on-line performance and statistical information, rather than hardcopy reports. They want to control what they receive and how often they receive it. Consistent with

much of the information gathered by the User Needs Assessment team, they would like to be able to access and print information at their discretion, reducing paper flow. They would like information about schedules, events, and project status that affects the way they do their jobs.

3.3 CCF User Services

The committee's starting point for quantifying information was through the experiences and logs of CCF User Services. User Services is chartered to support all CCF users' questions; however, most questions concern use of the CRAYs or VAXes. (Questions concerning system support and installation of IAS systems usually get passed to another group.) The matrix in Attachment 2 shows the following:

- Types of information requested on a monthly basis
- Methods for disseminating the information
- Potential number of users of the information
- Types of users reached
- Estimated size (in pages) of the information in hardcopy form
- Estimated annual percentage of increase (in pages)
- Importance of the information to users

The types of information detailed below correspond to the categories in the matrix. The matrix shows that about 7000 requests for information are filled on a monthly basis via phone, hardcopy, classes, or electronic mail. The priority column rates the importance of this information to the user, as defined by User Services.

I-II. Facility and Machine Information refers to the ACF facility alone; the UNICOS Userguide is the only place where this information is listed. No complete description exists of all the facilities that the CCF provides to Ames, such as the graphics facility, centrally administered machines, the Division's own Suns and SGIs, or the pass-through capability to reach the outside world via Pioneer.

Status Information is vitally important to the users, as evidenced by the number of monthly requests (3200). This is the most frequently requested type of information, and can only be distributed over the telephone. The electronic distribution listed in the matrix refers to scheduled status information, such as the machine being unavailable for preventive maintenance. Ironically, when the machine is down, users cannot access it to ascertain its status.

III. Events Notification of workshops and events are mass mailed and announced in *news* electronically.

IV. Training courses and material are available only through courses taught by the User Services staff and Lurnix. There is no C or C++ material, and none for any of the editors, except for vendor supplied documentation.

V. Alerts are distributed through the monitored mailing list to system administrators.

VI. Policies have been developed for the ACF but have not been distributed to the general user. No formal policies have been developed for the other systems.

VII. Procedures have been developed for the existing ACF policies. Again, these have not been distributed to the general user.

VIII. Operating systems - Information about the CRAY Y-MP operating system, is available only in the UNICOS Userguide. Information about many of the systems provided under the IAS contract is contained in the Workstation System Administration Guides I, II, and III. Some systems, such as the VAXes, are only covered by vendor manuals.

IX. Network Information is a time-critical component in a user environment. The only method to ascertain the status of local area networks is by calling User Services or the Integrated

Network Operations Center (INOC). No method exists for the user to get this information directly.

X. Periodicals are circulated by the CIC, but only within the Division.

XI-
XIII. Reports are locally and selectively distributed. There are often requests from a wider circle of users.

XIV. User Guides - After the initial releases (between 400 to 1000 copies), new requests for this material are still between 25 and 60 per month. These guides are only available in hardcopy.

XVI. Software - Most information about software is provided only through the vendor manuals.

4.0 Assessment of Current Services

The current methods for providing information, as well as the types of information provided, are appropriate and useful. However, it is clear that some methods are less successful than others and need augmentation. Outdated CIC functions could be eliminated with little loss, and useful functions could be incorporated into information services support. Some of the information that needs augmentation includes:

- Status information about the CRAY is unavailable when that system is down. This is exactly when users need status most. Users can call CCF User Services, but the lines are often busy when there are problems. It is clear that status information should be provided from another computer source.

- Events or notices that are distributed via *news* as part of the logon procedure, or are printed from computer output, have the same problem—if the computer is down, the information is not distributed. Additionally, users might find the information very important but are not using the computer at that time.

- Hardcopy user guides are involved and time-consuming to update. Changes in hardware or software take months to research, write, and distribute to users. Most users prefer a hardcopy manual; however, if the manual were maintained on line, users could access changes and print them locally.

- There is a general problem of gaining access to the information even when it is known to exist. There is no single place where users can look for information concerning the Division and the CCF.

5.0 Target Communities

The user community at Ames is a complex mixture of talents, skills, and disciplines. Its diversity, coupled with the fact that each person can wear many hats at the same time, makes the breadth of information—as well as the timely availability of accurate information—very important to the productivity of each individual. The types of information of interest to each person depend on what position they occupy on a number of levels, such as computer experience, job category, and area of research; additionally, their position in this multidimensional space changes over time.

Each person, therefore, may belong to many different "special interest groups," depending on which of the dimensions one chooses to focus. For example, an individual described as a first-time computer user whose job is to do research in material science, may benefit from getting information relevant to first-time users (such as training, facilities orientation, procedures, and policies), information of interest to all researchers, and information dealing specifically with material science. The types of information we have identified reflect that users' needs occur at multiple levels, over multiple dimensions. Examples of possible dimensions are shown in Attachment 3.

6.0 Recommendations

After assessing the information requirements of our user community in a short time frame, the committee makes the following recommendations.

6.1 General Recommendations

In order to give our customers the quality of service they require, it is imperative that the Division provide a centralized service for disseminating information, on a stable platform with a good uptime record. Although this will not answer all users' needs, it is the one solution that provides the most answers. It also provides the cornerstone for addressing other needs, such as that for self-paced training. This service will provide a place where status information can be located for other systems in the CCF. Manuals can be stored on line and updated with a minimum amount of effort. It will provide a place where users can find announcements and schedules, information about bugs and bug fixes, policies and procedures, reports, and alerts.

The types of information available on this system should broaden over time, and the system be able to increase the quantity of information stored, with staged and systematic implementations of additional features and storage.

6.2 Specific Recommendations

1. The system should be connected to the Ames network, and information should be available to all users, except those with isolated workstations. Any user, regardless of terminal type, should be able to access all information, with the exception of graphics. There are a number of users without smart terminals (some estimates are as high as 10%) who need information. This should not be interpreted to mean that the information should only be stored as flat text. Ms. Walsh, the head librarian at Ames, has stated that users will not use flat text unless there is no alternative. What it does mean is that the information may have to be stored in multiple formats.

2. The information should be stored in such a way that users can easily find what they want. The information will certainly be grouped into communities of interest. Because we have a diverse community of users (discussed in Section 4.0), the index should have multiple entries pointing to the same information. We should take deliberate steps to eliminate having to guess where the information is stored.

3. The system should be able to accommodate formatted text with graphics. A staged implementation should be developed that provides information to an increasing number of users, while adding facilities such as searching, help, and browsing. Special consideration should be given to ensure that the increased complexity is not detrimental to the system's ease of use. The ability to display graphics may not be implemented immediately but should be included in all plans.

4. Users should be able to both view and print documents; this ability will depend on the user's terminal and printer. Printing is a requirement for users who need updated manuals. Most users want hardcopies of this material.

5. The new service should augment existing services. Only after the new information service proves successful should current services be reduced or eliminated. This recommendation bridges the gap for some users who may not be able to access the new system.

6. Initially, the priority in which information is implemented on the system will be guided by the documentation that are on hand and facilities which are easy to implement, including:

   - a replacement for *on_line*
   - system status
   - events notification
   - cert alerts
   - ACF policies
   - trip reports

- weekly reports
- user guides

However, no choices should be made that preclude more advanced services. For example, a mechanism for providing users the ability to correspond with one another on selected research topics should be undertaken as soon as the primary service stabilizes.

7. Due to the complexity of creating an on-line information system, the committee recommends that a project team be formed to answer other important questions. The current committee will become an advisory board, working in conjunction with the team. It is also recommended that a support team, who will later implement the services, be formed at the outset in order to be involved in the entire process. At least one person should be assigned to participate in all three groups. The project team will identify or develop the items listed below.

   - A detailed list of attributes around which this system should be built.

   - A detailed plan for preparing the system and releasing it to the users. The plan should include the development of a system user manual, which will be released simultaneously with the system. As each section of the plan is completed, it can be presented to the committee for approval and then implemented.

   - A system for cataloging information. The system must be easy to use and the information readily accessible.

   - A method or criterion for choosing what information goes into the system.

   - The basic software that supports the system.

- Guidelines for data and system management, such as data ownership and update requirements.
- Tools that must be developed to maintain the system.
- Methods of informing the users about the system.
- Usage statistics on who is accessing the system and what kinds of information are being accessed.

7.0 Further Considerations

There are many elements to consider when deciding requirements for an on-line system. For example, the majority of users would like to be able to write, view, and print text at their desks without regard to the originating software, operating system, or destination. This is an impossible task in the heterogeneous computer environment at Ames. As a primary interface, a variety of platforms are available to each user, ranging from the CRAY Y-MP to Macintoshes, with workstations of all kinds in between. The display environment on desktops ranges from dumb terminals to sophisticated color graphic displays and workstations.

The diversity in markup languages ranges from typesetting software such as PageMaker and FrameMaker, to line editors such as emacs and vi, to camera-ready display languages such as TeX and troff. "Markup language" refers to additional information interspersed with the actual text of a document, which separates the document's logical elements and often specifies processing functions to be performed. It allows data to be stored, accessed, edited, published, and manipulated by specifying structural and procedural information, which is required by computer systems supporting various applications.

Lack of standards for transferring text and pagination markings from one environment to another compounds this problem. Most languages are able to output to a PostScript printer, which appears to be one constant feature. However, considering PostScript as a common element gives rise to other problems: PostScript previewers,

which exist on multiple platforms, are all slow and cpu intensive; documents are stored as images and cannot be searched; the display quality of fonts is poor (if unenhanced) due to the difference of 72 dpi vs 300 dpi for display fonts vs printer fonts. Certain fonts are unreadable on the screen without magnification, and some fonts are standardly available in some environments but not in others.

Today there is an emerging standard supported by DoD for a markup language called Standard Generalized Markup Language (SGML). However, because it is very early in the standardization process, interfaces and filters to printers and to other markup languages are not readily available on all platforms.

Attachment 1
Types of Information

- CCF Machine information
  - type
  - specifications
  - hardware
  - software
  - support
- Available facilities distributed by CCF
  - scientist workbench
  - AVS
  - Special software
- Machine status
- Mailing list subscriptions
- Events
  - workshops
  - lectures
  - parties
- Training events
- Training self-paced
  - UNICOS
  - NQS
  - languages
    - emacs
    - vi
    - getting started
    - C
    - C++
    - Fortran
    - dlib (distributed library)
  - graphics
  - Unix
  - vectorization
  - libraries
  - portability across platforms
- Security
  - information
  - password requirements
- Alerts
  - CERT
- Policies
- Procedures
  - how to get accounts
  - how to buy under the IAS contract
  - how to get documentation
- General CCF info
- Sun info
  - patches
- SGI info
  - patches
- DEC info
  - patches

Attachment 1
Types of Information

- CRAY info
  - vital statistics
  - local features
  - NQS
  - charges
  - procedures
  - accounts
  - access
  - queues
  - printing
  - allocation
  - scripts
  - scratch space
- News
  - down-time schedule
  - new versions installation
  - problems
- Newsletters
  - any newsletter which is of interest and can be available on line
- Graphics lab info
  - facilities
    - X Windows support
    - Solitaire (High Resolution Film Recording)
    - video animation system
    - 3-D software
    - SGI hardware
  - software supported
    - Plot3D. DISSPLA, NCAR Graphics, G.A.S., ARCGRAPH, GPLOT
- Network access info
  - dial-in access
  - connectivity requirement
- Periodical distribution
- CCF reprints
  - publications
  - technical reports
- Trip reports
- Weekly status reports
  - RC,RCU,RCA,RCS
  - Sterling ACF
- Documentation
  - User guides
    - system management guide
    - UNICOS Userguide
    - migration guide across multiple platforms
  - available for ordering
  - available for browsing
- Software
  - bug list
  - list of software available at ARC
  - compatibility across platforms
  - supported list
  - user supported list
  - unsupported list
  - documentation availability

Attachment 1
Types of Information

Output

- Central Print Facility
- HRFR
- VAS
- printers

	INFORMATION TYPE	ACCESS	MEDIA	# OF USERS	TYPE	PAGES	RATE	PRIORITY	COMMENTS
		Access/Month for information usually from User Services							
			The media used: P for phone, H for hard copy, E for electronic, C for training class						
				Potential number of users					
					R for researcher or technitial, M for manager				
						Estimation in pages			
							% Growth per year		
								Priority: H for High, M for medium, L for low	
									General comments
I	FACILITY INFORMATION								
	CSF	20	PHE	800	R	10	10	M	Cray Users
	CPF	30	PHE	1600	RM	10	10	M	
	GRAPHICS LABORAT	10	P	100	R	20	25	H	Advanced Researchers
	USER SERVICES	1	PH	1600	RM	1	0	M	
II	MACHINE INFORMATION								
	SPECIFICATION	10	PHE	100	RM	5	5	M	New Users
	STATUS	3200	PE	1000	RM	5	5	H	Cray, CCF Systems wknd schedule
III	EVENTS								
	WORKSHOPS	800	HE	1600	RM	20	20	L	Notices sent out
	LECTURES	800	HE	1600	RM	20	20	L	Notices sent out
IV	TRAINING								
	LURNIX	150	C	1000	RM	200	20	M	Wkstn Users and Admin (includes lab)
	UNICOS	30	CPH	800	R	300	25	M	Cray Users
	LANGUAGES								
	C	5	P	1000	R	100	10	M	Wkstn User
	FORTRAN	20	CPH	1200	R	150	5	M	Cray and Wkstn Users
	EDITORS								
	EMACS	1	P	520	R	20	0	M	Advanced Users
	VI	5	P	1200	R	30	0	M	All Users

INFORMATION TYPE	ACCESS	MEDIA	# OF USERS	TYPE	PAGES	RATE	PRIORITY	COMMENTS
	Access/Month for information usually from User Services							
		The media used: P for phone, H for hard copy, E for electronic, C for training class						
			Potential number of users					
				R for researcher or technitial, M for manager				
					Estimation in pages			
						% Growth per year		
							Priority: H for High, M for medium, L for low	
								General comments
V ALERTS								
CERT	250	E	250	RM	5	5	H	Sys Admin
VI POLICIES	10	P	800	RM	10	10	L	Current Cray Policies
VII PROCEDURES	0	-	800	RM	0	10	M	Cray Users
VIII OPERATING SYSTEMS								
UNICOS	250	P	800	RM	5	5	H	Cray Users
SUN/OS	100	P	500	R	5	10	H	Wkstn Users
VMS	100	P	500	R	5	15	M	VMS Users
IRIS	50	P	200	R	5	10	H	Iris Users
MPP	0	-	50	RM	0	25	H	None currently
IX NEWSLETTERS	0	-	2000	RM	400	0	H	Previous distribution
X NETWORK INFO	25	P	1200	RM	4	5	H	Any User
XI PERIODICALS	600	H	500	RM	60	10	M	
XII TRIP REPORTS	5	HE	100	RM	30	10	L	
XIII TECHNICAL REPORT	10	H	1200	RM	1000	0	M	

INFORMATION TYPE	ACCESS	MEDIA	# OF USERS	TYPE	PAGES	RATE	PRIORITY	COMMENTS
	Access/Month for information usually from User Services							
		The media used: P for phone, H for hard copy, E for electronic, C for training class						
			Potential number of users					
				R for researcher or technitial, M for manager				
					Estimation in pages			
						% Growth per year		
							Priority: H for High, M for medium, L for low	
								General comments
XIV WEEKLY REPORTS	30	H	150	M	30	0	L	
XV USER GUIDES								
WKSTN SYS ADMIN	25	H	1000	R	800	10	M	Admin and User (Avg 400/yr)
UNICOS	60	H	800	R	400	10	M	Cray Users (Avg 500/yr)
XV SOFTWARE								
BUG LIST	0	-	1200	R	0	25	H	None currently
CCF SUPPORTED	25	P	1200	R	10	10	M	
CCF UNSUPPORTED	10	P	1200	R	5	5	M	
USER SUPPORTED	10	P	1200	R	2	5	L	

Attachment 3
Identified User Communities

User's Experience:

- First-time users
  - help, training, facilities, policies, procedures
- Intermediate users
  - news, status, languages, compilers
- Advanced users
  - machine status, graphics, tools

Tools and Facilities Used by the Individual:

- Hardware Platforms
  - Cray, SGI, DEC, SUN, Intergraph, HP, Mac, PC, MPP
- Software
  - Operating systems : VMS, UNICOS, Unix, MPP
  - Layered products: C, Fortran, Macsyma, Gauss90, Nastran

Field of Research:

- Physics: Aerodynamics, Astronomy, Astrophysics, Atmospheric Mathematics
- Computer Science: Graphics, Molecular Modeling, Artificial Intelligence
- Chemistry: Material Science, Molecular Chemistry
- Biology: Molecular Biology
- Engineering: Structural, Mechanical, Acoustical, Simulation
- Psychology: Human Factors Analysis

Job Description:

- Support: Operations, system managers, system programmers, software specialists, hardware support, facilities, network support, administrative staff
- Management
- Researcher
- Developer

Others

- Resource monitor
- Colleague (peer-to-peer, non-CCF)
- Center-to-Center
- Remote connectivity

The second report was produced to answer questions David Fisher raised as a result of the first report.

August 28, 1992

Computer Information Services Project

In response to your memo dated June 29, 1992, the Information Services Committee has developed a high-level set of functional requirements (see attached) for electronic information distribution. Based on these requirements and a subsequent analysis, the committee recommends a two-phased approach for satisfying these functional requirements. We have identified specific actions, which are discussed below.

Through industry analysis, we have discovered that few vendors exist who provide a complete solution for RC's needs, and we believe that some important pieces of these solutions are still missing. Standards are still emerging, and vendors have targeted specific and understood areas for their development, resulting in a "patched" approach to implementing these standards. Vendors typically provide partial solutions, which can subsequently be integrated with another vendor's product. Additionally, HPCC, NASA, other ARC divisions, CCF, and NAS are actively exploring on-line development and information distribution requirements; therefore, choosing a specific vendor or selecting a standard at this time is premature. We believe that the recommended solution provides the flexibility we require in order to adopt information development and documentation standards as they emerge, while meeting the majority of our information distribution requirements now.

A prototype system (Gopher/WAIS, implemented cooperatively with NAS) now being tested, consists of a software package available in the public domain and installed on Pioneer. Users will be able to

access information and documentation from both NAS and CCF through a common viewer interface. However, as seen in the attached matrix, this system falls short of satisfying all the requirements, notably in the areas of standardization and graphics interfaces, and cannot be used as a long-term solution.

In Phase I, the computational capability contract will be tasked with making the prototype system fully operational and accessible to CCF users. This will require the creation of two plans: one for implementation of the on-line services and another for operation of these services. The implementation plan will describe the organization and procedures for data ownership, data maintenance, and system maintenance. The types of information to be provided are described in the May 22, 1992 committee report. The operational plan will include details for making the information available to users, including a schedule for announcing the service to CCF users.

Resources needed for Phase I equal two Full-Time Equivalents (FTEs) needed for three months. One FTE from the ACF staff is needed for operational and sustaining support. No additional hardware is anticipated in the short term; however, the present committee will address this issue four months after initial implementation, taking into account the response time requirement for user access, and the load on the present system. The current software is in the public domain, and we do not anticipate buying additional software for this phase. Currently, this software requires 40 MB of disk space, and disk partitions will have to be reconfigured to accommodate the future needs of the on-line facility. Some training will be required for the ACF FTE.

Phase II will focus on following the trends of this relatively new industry. The goal of this effort will be to decide at what point the technology and standards are stabilized enough to provide the functionality needed for CCF users, as well as positioning the CCF to take advantage of this new technology. The committee recommends assigning one or more individuals to track the technological progress and standardization activity of industry, academia, and the numerous government agencies. It is expected that as a result of this effort,

subsequent projects will be developed, which will adopt these standards as required.

Purpose: Information distribution system to serve all CCF users at Ames	Vendors	VERITY	INTERLEAF	PASSAGES	EL. BK TECH.	Gopher/Wais
Unix based centralized information system-client-server model		yes	yes	yes	yes	yes/no (metrics problem)
Accessible from X terminal with standard configuration(OSF Standards)		yes	yes	yes	yes	yes
Accessible from standard dumb terminals		yes-seperate product	yes	no	no	yes
Utilizing a minimum number of vendors (single vendor solution preferred)		yes	yes	yes-multiple vendors integrated.	2 vendors	free
User Interface						
	Ability to input documents from a variety of authoring environments	yes-filters	yes	yes	yes-Avalanche	no
	User friendly	yes	yes	yes	yes	yes
	Good appearance	yes	yes	yes	yes	yes
	Easy for first time use	yes	yes	yes	yes	yes
Consistency between viewed document and printed document		no	yes	yes	yes	yes
System	Navigeable with Contextual Reference	no	yes	yes	yes	no
	Easy to maintain and modify	no	no	no	no	yes
	Seperate viewing and production system from data-document maintained in there acquired forma	yes	no	yes	yes	yes
	Not labor or resource intensive	no	no	no	no	yes
	Ability to integrate graphics,video,sounds,color ,hyperlink in the text.	hyperlink only	graphics only	yes	yes	sound only
	Data should be easily maintainable	no	no	no	no	yes
	Expandable	yes	yes	yes	yes	yes
	Ability to restrict file view on a per user basis	yes	yes-additional product	yes	yes	no
User tools	Searching-within documents -across documents	yes	yes	yes	yes	yes
	Journaling-public and private	yes	yes	yes	yes	yes
	Annotation-public and private	full only	full only	yes	yes	full only
	Printing -full file or selected part	yes	yes	yes	yes	no
	Output	yes	yes	yes	yes	yes
	Bookmarking	no	yes	yes	yes	yes
Emerging standards have to be fully implementable (CALS and SGML)		no	yes-non native	yes	yes	no
Platform size for production system(2000 pages)		.5hr Sparc-22mips	1hr Sparc-40 mips	.5-1hr Sparc-40 mips	.5-1hr Sparc-40 mips	20 min. est.
Platform size for viewing system with user response no greater than 5 sec.		small	small	small	small	small
Metrics	User access	yes	yes-additional sotware	no-custom software	no-custom software	yes-dependent on
	File access	no/yes	yes-additional sotware	no-custom software	no-custom software	yes-dependent on

REQUIREMENTS

General requirements were obtained by reviewing available literature, analyzing current offerings from the commercial sector, and holding in-depth discussions with other supercomputing centers. We researched different architectural design and functional characteristics for standard information access, distribution, and storage requirements.

Specific local requirements were gathered by interviewing various members of the Division, and incorporated user-related findings derived from the User Needs Assessment team.

The technology and standards for on-line information services are rapidly changing. We can, therefore, expect greater emphasis and significant change in this discipline as it matures over the next two years. To illustrate this, a group representing all the NFS supercomputing centers, supercomputer and workstation vendors, and a technical publisher have been meeting annually for several years. The participants discuss the requirements for standardization and sharing of on-line documentation and information distribution. Their findings have been published, are considered appropriate for our environment, and have been incorporated with requirements.

Gantt Chart Report
inf.gantt

As of Date: 08/22/91